WIPO WORLDWIDE SYMPOSIUM ON THE INTELLECTUAL PROPERTY ASPECTS OF ARTIFICIAL INTELLIGENCE

Stanford University,
Stanford (California), United States of America

March 25 to 27, 1991

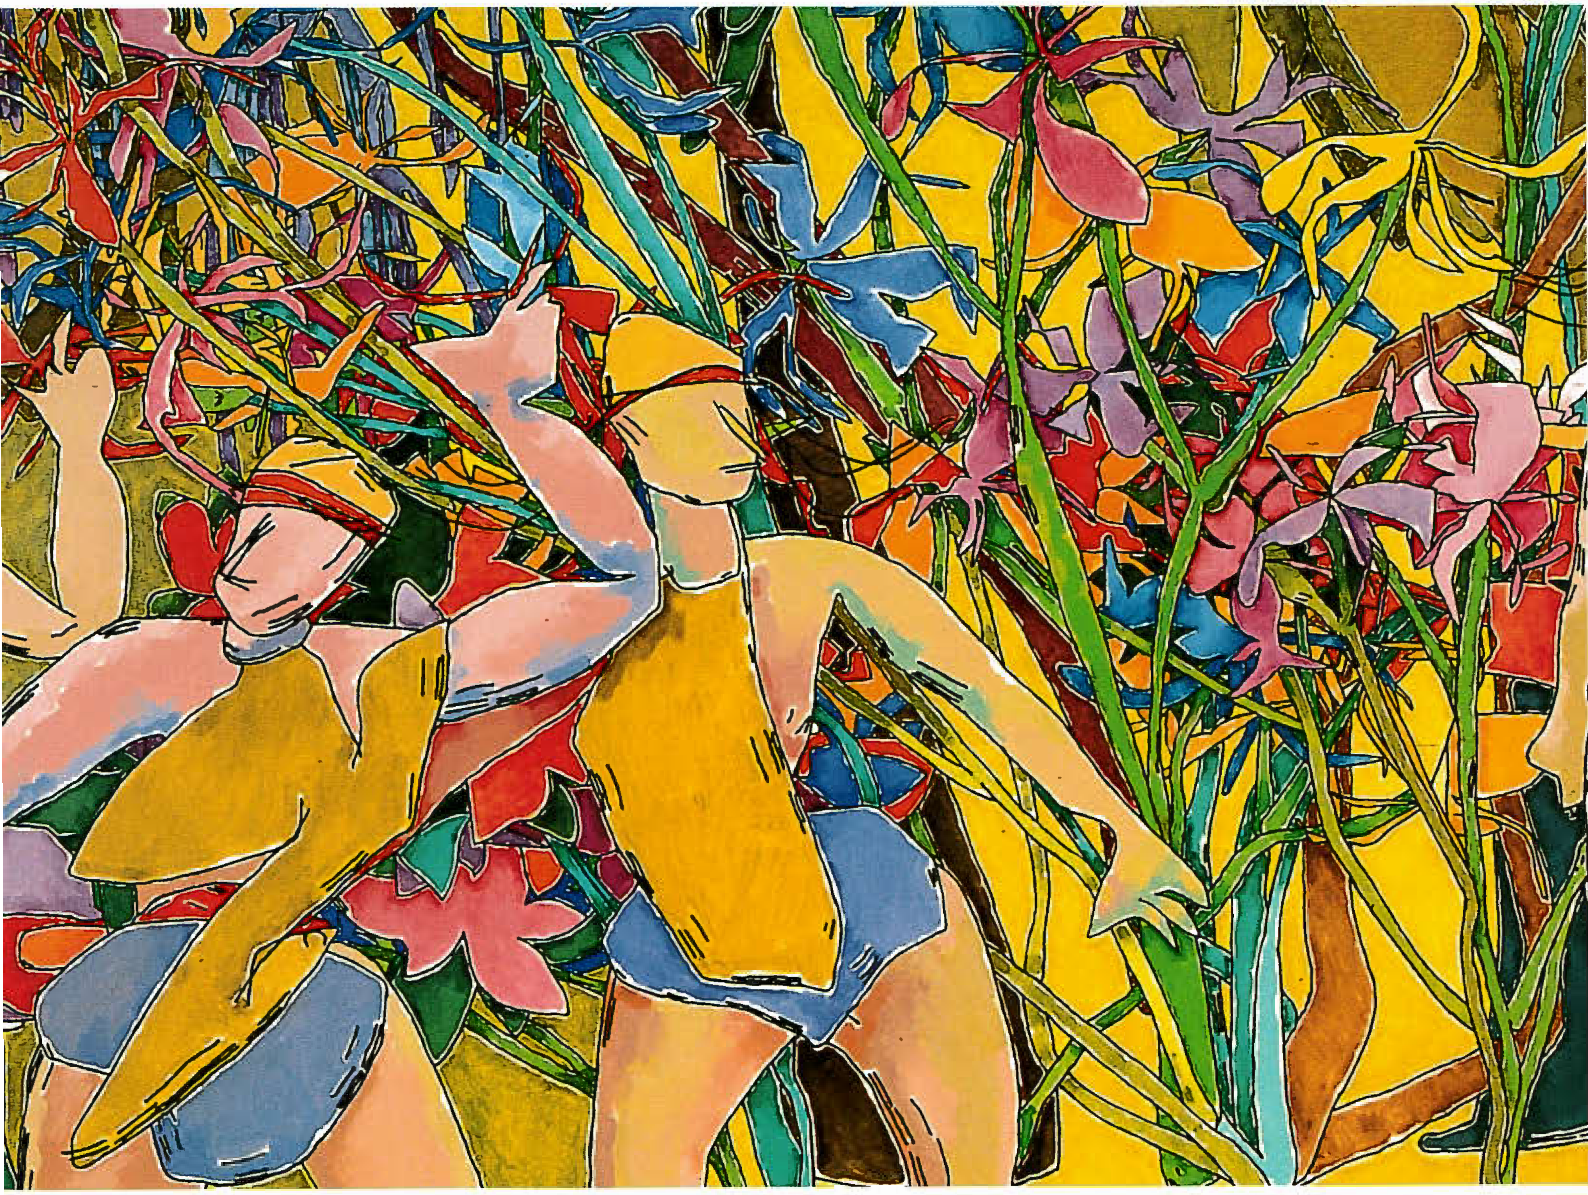

World Intellectual Property Organization

WIPO WORLDWIDE SYMPOSIUM ON THE INTELLECTUAL PROPERTY ASPECTS OF ARTIFICIAL INTELLIGENCE

Stanford University,
Stanford (California), United States of America

March 25 to 27, 1991

World Intellectual Property Organization

GENEVA, 1991

The cover illustration, "Coming Into a Lighter Place",
was produced by the AARON computer program created
by Dr. Harold Cohen
of the University of California at San Diego, United States of America.
It appears here with the gracious consent of Dr. Cohen.
The image was photographed by Becky Cohen and is in the collection
of Robert and Deborah Hendel.

WIPO PUBLICATION No 698 (E)
ISBN 92-805-0367-7

WIPO 1991
Reprinted 1995

PREFACE

The World Intellectual Property Organization (WIPO) organized a Worldwide Symposium on the Intellectual Property Aspects of Artificial Intellegence at Stanford University, Stanford, California, United States of America, from March 25 to 27, 1991.

"Artificial intelligence" is an expression commonly used to designate those types of computer systems that display certain capabilities associated with human intelligence, such as perception, understanding, learning, reasoning and problem-solving. The three best-known existing categories of artificial intelligence are: expert (or knowledge-base) systems, perception systems and natural language systems.

These three "systems" are the results of state-of-the-art research, development and application in the field of artificial intelligence. A longer-term objective is envisaged, however, which is to combine the various specific "systems" into ever more comprehensive universal systems which--being able to perceive, to understand, to learn, to communicate, to reason and, most important from the standpoint of intellectual property, also to create--will perform certain functions of human intelligence with ever increasing accuracy.

The Symposium examined the various categories of artificial intelligence and their main fields of application in the context of their possible intellectual property implications. Discussions focused on the form of protection that should apply to certain existing creations protected by intellectual property rights in respect of their inclusion in artificial intelligence systems (e.g., in expert systems), on what protection, if any, should be granted to artificial intelligence systems themselves, and on determining the appropriate intellectual property status of the outputs of such systems (such as computer-produced creations).

On the basis of this examination, an attempt was made to assess the possible long-term impact of artificial intelligence on intellectual property law (in light of the fact that existing international systems for the protection of intellectual property have been established with human creations in mind and that, with the advent of artificial intelligence, the possibility of "artificial creation" is emerging).

The Symposium included presentations by experts from the fields of computer science, law (both professors and practitioners) and business. The present volume contains the texts of those presentations.

The audience consisted of some 80 persons, including government officials, representatives of international and international non-governmental organizations, and the general public. Their list is included in this volume.

The World Intellectual Property Organization is grateful to Stanford University for hosting the meeting. It expresses its thanks to all the speakers, and in particular to Professor Paul Goldstein, whose advice in respect of the organization of the meeting was of the greatest importance.

A. Bogsch

Arpad Bogsch
Director General
World Intellectual Property Organization

August 1991

TABLE OF CONTENTS

page

PREFACE i

TABLE OF CONTENTS iii

PROGRAM 1

INFORMATION ON THE SPEAKERS 5

Opening Statement 17

Dr. Arpad Bogsch
Director General of WIPO

Welcoming Remarks 19

Mr. Ralph Oman
Register of Copyrights
Copyright Office
United States of America

Welcoming Remarks 21

Dr. Donald Kennedy
President of Stanford University

I. Notion and General Overview of Artificial Intelligence

Notion and General Overview of Artificial Intelligence 23

Professor Terry Winograd
Stanford University
Palo Alto, California
United States of America

Artificial Intelligence and Intellectual Property Law:
Some Problematical Examples 31

Professor L. Thorne McCarty
Rutgers University
New Brunswick, New Jersey
United States of America

The Machine Translation System and Its Possible
Intellectual Property Implications 39

Makoto Nagao
Kyoto University
Kyoto, Japan

II. Main Categories of Artificial Intelligence and their Intellectual Property Aspects

Main Categories of Artificial Intelligence and
their Intellectual Property Aspects 45

Mr. Andy Johnson-Laird
Johnson-Laird, Inc.,
Portland, Oregon
United States of America

Copyright Protection for Artificial Intelligence Systems 55

Morton David Goldberg
Schwab, Goldberg, Price and Dannay
New York, N.Y.
United States of America

Protecting Expert Systems, in Particular Expert System Knowledge:
a Challenge for Lawyers 77

Jaap H. Spoor
Vrije Universiteit and Trenité van Doorne
Amsterdam, Netherlands

The Ontologies of Knowledge Representations:
Why Knowledge Definitions Should Not Be Protected 85

Daniel M. Russell
Systems Sciences Laboratory
Xerox PARC Center
Palo Alto, California
United States of America

IA and IP: Intelligent Applications and
Intellectual Property 95

Stephen A. Weyer
Manager, Intelligent Applications
Advanced Technology Group
Apple Computer, Inc.
Cupertino, California
United States of America

III. Main Categories of Artificial Intelligence and their Intellectual Property Aspects (continuation)

Intellectual Property and Software:
the Assumptions are Broken 101

Randall Davis
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts
United States of America

The Patentability of Artificial Intelligence Under U.S. Law 121

Ronald S. Laurie
Irell & Manella
Menlo Park, California
United States of America

Intellectual Property Law Aspects of
Artificial Intelligence 151

Thomas K. Dreier
Staff Member
Max-Planck-Institute for Foreign and International
Patent, Copyright, and Competition Law
Munich, Germany

Intellectual Property Aspects of Artificial
Intelligence: Some Issues 169

Ashok Bhojwani
Managing Director
TSG Consultants Limited
New Delhi, India

Main Categories of Artificial Intelligence and their
Intellectual property Aspects 189

Clifford M. Sturt
Partner
Marks & Clerk
London
United Kingdom

IV. Computer-Produced Creations

Computer Produced Creations 201

Rudolf v. B. Rucker
San Jose State University and Autodesk, Inc.,
Sausalito, California
United States of America

Proprietary Rights in Computer-Generated Productions 205

J. H. Reichman
Professor of Law
Vanderbilt University
Nashville, Tennessee
United States of America

Computer-Produced Creations 225

Robert Barr
Partner, Townsend and Townsend
Palo Alto, California
United States of America

Copyright Protection of Computer-Produced Creations 229

Péter Gyertyánfy
Hungarian Bureau for the Protection of Authors' Rights
(ARTISJUS)
Budapest, Hungary

Computer-Produced Creations 235

Akin Thomas
Managing Director
Heinemann Educational Books
Ibadan, Nigeria

V. <u>The Impact of Artificial Intelligence on the Law of Intellectual Property</u>

Computers and Authorship: The Copyrightability
of Computer-Generated Works 241

Arthur R. Miller
Bruce Bromley Professor of Law
Harvard Law School
Cambridge, Massachusetts
United States of America

Artificial Intelligence in the Film Industry: an Oxymoron? 271

Norman Alterman
Senior Vice President, Legal Affairs
Motion Picture Export Association of America (MPEAA)
New York, N.Y.
United States of America

Intellectual Property Problems Related to
Artificial Intelligence Systems 279

Antonio Millé
Latin-American Institute for Advanced Technology
Computer Science and Law
Buenos Aires, Argentina

The Creation of Works of Copyright Under
Japanese Copyright Law Resulting From the Utilization
of Artificial Intelligence 293

Shigeru Miki
Miki Law & Patent Office
Tokyo, Japan

General Report 297

Paul Goldstein
Stella W. and Ira S. Lillick Professor Law
Stanford University
Palo Alto, California
United States of America

LIST OF PARTICIPANTS 301

PROGRAM

Monday, March 25, 1991

Opening Ceremony

Opening Statement: Arpad Bogsch, Director General of WIPO

Welcoming Remarks: Ralph Oman, Register of Copyrights, Copyright Office, United States of America

Welcoming Remarks: Donald Kennedy, President of Stanford University

I. Notion and General Overview of Artificial Intelligence

Moderator: Arpad Bogsch, Director General of WIPO

Speakers: Terry Winograd, Stanford University, Palo Alto, California, United States of America

L. Thorne McCarty, Rutgers University, New Brunswick, New Jersey, United States of America

Makoto Nagao, Kyoto University, Kyoto, Japan

Discussion

II. Main Categories of Artificial Intelligence and their Intellectual Property Aspects

Moderator: Ralph Oman, Register of Copyrights, Copyright Office, United States of America

Speakers: Andy Johnson-Laird, Johnson-Laird, Inc., Portland, Oregon, United States of America

Morton David Goldberg, Schwab, Goldberg, Price and Dannay, New York, N.Y., United States of America

Jaap H. Spoor, Vrije Universiteit and Trenité van Doorne, Amsterdam, Netherlands

Daniel M. Russell, Systems Sciences Laboratory, Xerox PARC Center, Palo Alto, California, United States of America

Stephen A. Weyer, Manager, Intelligent Applications, Advanced Technology Group, Apple Computer, Inc., Cupertino, California, United States of America

Discussion

Tuesday, March 26, 1991

III. Main Categories of Artificial Intelligence and their Intellectual Property Aspects (continuation)

Moderator: Mihály Ficsor, Director, Copyright Law Division, WIPO

Speakers: Randall Davis, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America

Ronald S. Laurie, Irell & Manella, Menlo Park, California, United States of America

Thomas K. Dreier, Staff Member, Max-Planck-Institute for Foreign and International, Patent, Copyright, and Competition Law, Munich, Germany

Ashok Bhojwani, Managing Director, TSG Consultants Limited, New Delhi, India

Clifford M. Sturt, Partner, Marks & Clerk, London, United Kingdom

Discussion

IV. Computer-Produced Creations

Moderator: Francis Gurry, Special Assistant, Office of the Director General, WIPO

Speakers: Rudolf v. B. Rucker, San Jose State University and Autodesk, Inc., Sausalito, California, United States of America

Jerome H. Reichman, Professor of Law, Vanderbilt University, Nashville, Tennessee, United States of America

Robert Barr, Partner, Townsend and Townsend, Palo Alto, California, United States of America

Péter Gyertyánfy, Hungarian Bureau for the Protection of Authors' Rights, (ARTISJUS), Budapest, Hungary

Akin Thomas, Managing Director, Heinemann Educational Books , Ibadan, Nigeria

Discussion

Wednesday, March 27, 1991

V. The Impact of Artificial Intelligence on the Law of Intellectual Property

Moderator: Arpad Bogsch, Director General of WIPO

Speakers: Arthur R. Miller, Bruce Bromley Professor of Law, Harvard Law School, Cambridge, Massachusetts, United States of America

Norman Alterman, Senior Vice President, Legal Affairs, Motion Picture Export Association of America (MPEAA), New York, N.Y., United States of America

Antonio Millé, Latin-American Institute for Advanced Technology, Computer Science and Law, Buenos Aires, Argentina

Shigeru Miki, Miki Law & Patent Office, Tokyo, Japan

Paul Goldstein, Stella W. and Ira S. Lillick Professor Law, Stanford University, Palo Alto, California, United States of America

Discussion

Closing of the Symposium by the Director General

- * -

INFORMATION ON THE SPEAKERS

Norman ALTERMAN

Mr. Norman Alterman, a national of the United States of America, is Senior Vice President for Legal Affairs of the Motion Picture Export Association of America (MPEAA), in New York. He is also Legal Adviser to the International Federation of Film Producers Association (FIAPF).

Mr. Alterman was a member of the Ad Hoc Working Group on United States Adherence to the Berne Convention, convened by the U.S. Department of State. He has lectured widely on various copyright topics. Mr. Alterman holds several professional memberships, including the Copyright Society of the U.S.A., the American and New York State Bar Associations, and the International Bar Association.

Mr. Alterman was graduated from the Wharton Business School, and received his J.D. from Harvard Law School.

Robert BARR

Mr. Robert Barr, a national of the United States of America, is a partner in the Palo Alto office of the intellectual property law firm of Townsend and Townsend. He specializes in computer law, including patent prosecution, licensing, litigation, and software protection.

Mr. Barr was Manager of Product Development at Tesseract Corporation, a San Francisco software company, and has taught legal research and writing at the University of California Law School (Boalt Hall).

Mr. Barr earned a B.S. degree in Electrical Engineering from the Massachusettes Institute of Technology in 1970 and a J.D. from Boston University School of Law in 1973.

Ashok BHOJWANI

Mr. Ashok Bhojwani, a national of India, is Managing Director of TSG Consultants (P) Limited, New Delhi, India.

Mr. Bhojwani has had experience in the computer field since 1964, including 14 years experience with IBM in the United States and in India. His experience includes hardware design of super computers in IBM's "Advanced Computer Systems" laboratory, software development, technical and general management, senior executive education and information systems strategic planning.

Between 1979 and 1980, Mr. Bhojwani was a member of the Executive Council of the Computer Society of India. In 1986, he was Chairman of the Software Committee of the Manufacturers Association of Information Technology (MAIT), and between 1987 and 1988, he was Vice President of MAIT.

Mr. Bhojwani holds an M.S. (Electrical Engineering) from the Polytechnics Institute of Brooklyn, New York, and a B.Tech. (Hons.) (Electronics), from I.I.T., in Kharagpur.

Mr. Bhojwani has written and lectured extensively on subjects pertaining to the legal protection of computer software, including recent seminars sponsored by WIPO in several countries.

Randall DAVIS

Randall Davis, a national of the United States of America, is a Professor at the Massachusetts Institute of Technology in both the Department of Computer Science and the School of Management as well as Associate Director of the Artificial Intelligence Laboratory. He has been involved in issues of software and intellectual property for several years. In 1989 he served in a workshop on the issue organized by the National Research Council of the National Academy of Science, the results of which have recently been published as Intellectual Property Issues in Software, National Academy Press, 1991. He has spoken and organized panel discussions at MIT and elsewhere, and served as an expert witness in several software cases. In 1990 he served as special master to the court in a software copyright infringement case in the Second Circuit in New York.

Thomas K. DREIER

Mr. Thomas K. Dreier, a national of Germany, is on the staff of the Max-Planck Institute for Foreign and International Patent, Copyright and Competition Law in Munich, Germany. He is also a legal expert on copyright questions of cable and satellites for the Commission of the European Communities.

Mr. Dreier has written and lectured on topics related to copyright law, primarily computer programs and integrated circuits, cable and satellite programs, and the harmonization of copyright laws in the European Communities.

Mr. Dreier holds law degrees from the University of Munich, the University of Geneva and the New York University. He is a member of the New York State Bar, the International Bar Association, the International Association for the Protection of Literary and Artistic Property (ALAI), and the German Computer Law Association.

Morton D. GOLDBERG

Mr. Morton D. Goldberg, a national of the United States of America, is a partner at Schwab, Goldberg, Price and Dannay in New York. His practice focuses on copyright, with a particular emphasis on computer programs.

Mr. Goldberg has written and lectured extensively on various copyright-related topics. He was an adviser to the U.S. Copyright Office during the long process of revising the 1909 Copyright Law, and continues to play a role as a private-sector adviser in the formulation of the U.S. Government's copyright policy, particularly in its international aspects.

From 1985 to 1987, he served on the Ad Hoc Working Group on United States Adherence to the Berne Convention convened by the U.S. Department of State. He also served on the Advisory Panel for Information Technology and Intellectual Property of the U.S. Congress' Office of Technology Assessment.

Mr. Goldberg's professional memberships have included service as an Officer or Director of the American Intellectual Property Law Association, The Copyright Society of the U.S.A., The Computer Law Association, and The Patent, Trademark and Copyright Section of the American Bar Association.

Mr. Goldberg was graduated magna cum laude from Harvard University in 1951, and received his LL.B. from Yale University in 1954.

Paul GOLDSTEIN

Professor Paul Goldstein, a national of the United States of America, is the Stella W. and Ira S. Lillick Professor of Law at Stanford University. Professor Goldstein is the author of a three-volume treatise, "Copyright: Principles, Law and Practice" (Little, Brown and Company, 1989) a leading law school text on intellectual property, and numerous articles on copyright law.

Professor Goldstein has served as chairman of the U.S. Office of Technology Assessment Advisory Panel on Intellectual Property Rights in an Age of Electronics and Information, and has been a Visiting Scholar at the Max-Planck Institute for Foreign and International Patent, Copyright and Competition Law in Munich, Germany.

Professor Goldstein is a Trustee of the Copyright Society of the U.S.A. and is of Counsel to the law firm of Morrison & Foerster.

Péter GYERTYÁNFY

Mr. Péter Gyertyánfy, a national of Hungary, was appointed Legal Director of the Hungarian Bureau for the Protection of Authors' Rights (ARTISJUS) Budapest, on December 1, 1990. Previously, he worked for ARTISJUS as Deputy Chief of the Legal Department and, from 1985 to 1990, as Director of Administration.

Since 1983, Mr. Gyertyánfy has been temporary lecturer on international copyright law at the Law School of the Eötvös Loránd University, Budapest. In 1990, he was appointed Titular Assistant Professor of the University. That year, he earned his academic title by completing a thesis entitled "Copyright Law Protection of Computer Programs and Electronic Data Basis".

Between 1968 and 1979, Mr. Gyertyánfy held several positions in the banking and financial sector, including counselor to the International Financial Department of the Hungarian Ministry of Finance.

Mr. Gyertyánfy was graduated from the Law School of the Eötvös Loránd University in 1968.

Andy JOHNSON-LAIRD

Mr. Andy Johnson-Laird, a British national resident in United States of America, is President of Johnson-Laird Inc., a consulting company specializing in Techno-Archeology℠, the forensic analysis of software for litigation purposes. This analysis includes plagiarism assessment and failure analysis of software development projects.

Johnson-Laird Inc. has provided expert consulting services to major law firms in San Francisco, Los Angeles, Minneapolis, London, Brussels and Tokyo.

Mr. Johnson-Laird has 30 years experience of programming, managing programmers, and developing interoperable and competitive products on main-frames, minicomputers and microcomputers. He has authored several papers explaining software development technology to the legal community, especially in the context of new developments such as neural networks, and addressing the the issues surrounding the reverse engineering of software. His professional career has taken him from England to Portland, Oregon, via France and Canada.

Ronald S. LAURIE

Mr. Ronald S. Laurie, a national of the United States of America, is an attorney-at-law and partner at Irell and Manella, a law firm in Menlo Park, California.

Mr. Laurie worked as a computer systems designer and programmer in the aerospace industry for several years. He has been a practicing attorney since 1968. In August 1988, he joined Irell and Manella, a firm with headquarters in Los Angeles, California. As a resident partner in the Menlo Park (Silicon Valley) office of that firm, he heads its computer law group. He specializes in litigation and counselling regarding proprietary rights in computer technology, including software and semiconductor chip designs. A substantial part of his practice involves licensing and technology transfer with an emphasis on international transactions.

Mr. Laurie graduated with a B.S. degree in engineering from the University of California at Berkeley and holds the degree of J.D. from the University of San Francisco Law School. He is the author of several articles on the law on intellectual property rights in computer technology and is a contributing editor to the journal "Computer Lawyer."

L. Thorne McCARTY

Professor L. Thorne McCarty, a national of the United States of America, is on the faculty of the Department of Computer Science and the School of Law at Rutgers University in New Jersey.

Professor McCarty has written and lectured widely on topics involving computers, information systems and artificial intelligence, especially as they relate to the legal profession. He is a co-editor of the Massachusetts Institute of Technology Press series "Artificial Intelligence and Legal Reasoning," is on the editorial board of "Artificial Intelligence and Law: An International Journal" and is the reviewer for various journals on artificial intelligence and computers, and for several major book publishers.

Professor McCarty was graduated with honors from Yale University and received his J.D. from Harvard Law School.

Professor Miller has made frequent television appearances as a law commentator. He has also written occasional columns on legal subjects for various newspapers.

Professor Miller has held a number of public service positions in the fields of privacy, computers, copyright and courts, among them as a Commissioner on the United States Commission on New Technological Uses of Copyrighted Works (CONTU), as the Reporter for and member of the Advisory Committee on Civil Rules for the United States Supreme Court, and as the Reporter for the American Law Institute's Project on Complex Litigation.

Professor Miller has an undergraduate degree from the University of Rochester and a law degree from Harvard Law School.

Makato NAGAO

Professor Makato Nagao, a national of Japan, has been a professor of Electrical Engineering at Kyoto University since 1973. He has conducted research activities in the area of artificial intelligence for many years. Among his subjects of inquiry have been pattern recognition, image processing, natural language processing and machine translation.

Professor Nagao has participated in many international conferences and symposia on artificial intelligence. He is a contributing editor of several journals, including "Computer Vision," "Graphics and Image Processing" and "Artificial Intelligence." He has also written extensively in both English and Japanese.

Professor Nagao was graduated from Kyoto University in 1959 with a degree in Electrical Engineering, and received M.S. and Ph.D. degrees from the same institution in 1961 and 1966, respectively.

Jerome H. REICHMAN

Professor Jerome H. Reichman, a national of the United States of America, is a Professor of Law at Vanderbilt University in Nashville, Tennessee. He has held teaching positions at the University of Florida and at the Ohio State University.

Between 1975 and 1979, Professor Reichman served as an official of the International Trade Centre, UNCTAD/GATT, in Geneva. He was a founding member of the American branch of the International Association for the Protection of Literary and Artistic Property (ALAI), and has been active in the International Association for the Advancement of Teaching and Research in Intellectual Property (ATRIP).

Professor Reichman has written and lectured extensively on topics involving intellectual property with a particular focus on copyright protection for computer programs.

Professor Reichman was awarded a B.A. from the University of Chicago in 1955, and received his J.D. degree from Yale University in 1979.

Shigeru MIKI

Mr. Shigeru Miki, a national of Japan, is a senior partner in the firm of Miki Law And Patent Office in Tokyo. He is a specialist in patent, trade mark and copyright law, with an emphasis on computers.

Mr. Miki has written and lectured extensively on the legal protection of computer programs both in Japan and in other countries. Since 1988, he has been a member of the Research Committee on the Legal Protection of Computer Software sponsored by the Japanese Agency of Cultural Affairs. He is also a lecturer at the Chuo University.

Mr. Miki received his basic law degree from the Chuo University in 1968. In 1980, he was awarded an LL.M. in Trade Regulation from the New York University School of Law.

Antonio MILLE

Mr. Antonio Millé, a national of Argentina, is the Senior Partner of Estudio Millé, a law firm in Buenos Aires specializing in intellectual property matters, including copyright. He is counsel to the Argentine collecting society AADI-CAPIF, the Argentine software vendors' association CES, and the hardware manufacturers' association CICOM. He is also the executive secretary of the Latin American Federation of Performing Artists (FLAIE). Mr. Millé has written and lectured widely on copyright topics, with a special emphasis on legal protection for computer programs and other information related technology.

Mr. Millé holds many professional memberships, including the International Literary and Artistic Association (ALAI), the Computer Law Association, the Interamerican Copyright Institute (IDA) and the Licensing Executives Society (LES). He holds editorial positions on the journal "DAT - Derecho de Alta Tecnologia," "International Computer Legal Adviser," and "Software Protection."

Arthur R. MILLER

Professor Arthur R. Miller, a national of the United States of America, is the Bruce Bromley Professor of Law at Harvard Law School, where he has taught since 1971.

Before joining the Harvard faculty, Professor Miller practiced law in New York City and taught at the University of Minnesota and the University of Michigan. Professionally, he is known for his work on court procedure--a subject on which he has authored or co-authored more than 25 books--, copyright, unfair competition, and remedies. He has also written, testified on, debated, and helped formulate legislation on the right of privacy. His book "The Assault on Privacy: Computers, Data Banks, and Dossiers" (1971) reached a wide readership.

Rudolf v.B. RUCKER

Professor Rudolf Rucker, a national of the United States of America, is best known for his popular books on mathematics and for his science fiction novels, including the cyberpunk classics "Software" and "Wetware," each of which won the Philip K. Dick Award. Professor Rucker is Professor of Mathematics and Computer Science at San Jose State University, and Mathenaut in the Advanced Technology division of Autodesk Inc.

Professor Rucker has recently published his thirteenth book, "The Hollow Earth," a science-fiction adventure novel that starts in 1836 and features Edgar Allan Poe. He has also co-authored two science series software packages: "CA Lab" and "James Gleick's Chaos: The Software."

Professor Rucker graduated from Swarthmore College in 1967 and received a Ph.D. in Mathematics from Rutgers University in 1972.

Daniel M. RUSSELL

Mr. Daniel M. Russell, a national of the United States of America, is a Member of the Research Staff at Xerox Palo Alto Research Center (PARC). Currently, he studies issues of semi-formal knowledge representation, information visualization and agent-directed knowledge acquisition. From 1984 through 1991, he led the "Instructional Design Environment" project to develop a practical computer-aided design and analysis system for use in performing ill-structered design tasks. In addition to PARC, he is an associate member of the Institute for Research on Learning (IRL) in Palo Alto, and is an adjunct lecturer on the Engineering and Computer Science (Computer Science) faculty of the University of Santa Clara.

Mr. Russell's research interests involve knowledge representation, building and using systems to solve difficult design problems, and automatic problem-solving and planning systems.

Mr. Russell received his B.S. in Information and Computer Science from U.C. Irvine in 1977, and his M.S. and Ph.D. degrees in Computer Science from the University of Rochester in 1979 and 1984, respectively. While at Rochester, he did graduate work in the neuropsychology of laterality and computer vision.

Jaap H. SPOOR

Professor Jaap H. Spoor, a national of the Netherlands, is Professor of Intellectual Property Law, Vrije Universiteit, and Attorney, member of the Netherlands Bar, Trenité Van Doorne, Amsterdam. He specializes in copyright, computer law, trademark law and other intellectual property law topics.

Mr. Spoor also is a Board member of the Netherlands' Association for Copyright Law, Editor of the monthly magazine Informatierecht (Information Law), and Executive Committee member of the International Literary and Artistic Association (ALAI). He co-authored the leading Dutch treatise on copyright law, and has published several other books and numerous articles concerning copyright, industrial property and computer law, some of which deal with expert systems.

Clifford M. STURT

Mr. Clifford M. Sturt, a national of the United Kingdom, is a partner in the London firm of Marks and Clerk, which he joined in 1984.

Mr. Sturt is a specialist in intellectual property protection for computer-related technology. He has lectured widely on this subject, and has acted as an expert witness in disputes involving software copyright.

Mr. Sturt was graduated from the Imperial College of Science and Technology (London University) in 1977 with an honors degree in physics. He was also awarded an Associateship of the Royal College of Science.

Mr. Sturt qualified as a Chartered Patent Agent and European Patent Attorney in 1981.

Akin THOMAS

Mr. Akin Thomas, a national of Nigeria, has been managing director of Heinemann Educational Books - Nigeria since 1984.

Mr. Thomas has written on many copyright-related topics, with a focus on issues affecting publishers. His interest in copyright dates from his tenure as president of the Nigerian Publishers Association in 1983.

Mr. Thomas is active in several professional associations, including the Nigerian Publishers Association and the Nigerian Copyright Council.

Stephen WEYER

Mr. Stephen Weyer, a national of the United States of America, is Manager of Intelligent Applications in Apple Computer's Advanced Technology Group, Cupertino, California. Apple's IA group is investigating the advanced software development and learning environments, tools for knowledge representation, information access and conversational user interfaces based especially on artificial intelligence (AI) technologies that will enable next-generation applications.

Mr. Weyer has been with Apple for more than four and a half years; before that, he managed AI programming environments at Hewlett-Packard Laboratories and investigated educational hypermedia (electronic book) applications at Atari Research and at Xerox Palo Alto Research Center.

Mr. Weyer has a Ph.D. from Stanford University.

Terry WINOGRAD

Professor Terry Winograd, a national of the United States of America, is Professor of Computer Science at Stanford University.

Professor Winograd's research on natural language understanding by computers was the basis for his book "Understanding Natural Language," and his textbook "Language as a Cognitive Process," as well as for numerous articles in both scholarly journals and popular magazines. His most recent book, co-authored with Fernando Flores, takes a critical look at work in artificial intelligence and presents an alternative theory of language and thought, which suggests new directions for the design of intelligent human/computer systems. The book is entitled "Understanding Computers and Cognition: A New Foundation for Design."

Professor Winograd is on the editorial board of a number of journals, including "Artificial Intelligence," "AI Expert," "AI & Society," "Journal of Computing and Society," "Human-Computer Interaction."

Professor Winograd received his B.S. in mathematics from the Colorado College in 1966, and a Ph.D. in Applied Mathematics from the Massachusetts Institute of Technology (MIT) in 1970. He taught at MIT from 1970 to 1973, and has been on the faculty of the Computer Science Department at Stanford University since 1973. He also has appointments in the Department of Linguistics and the Program in Values, Technology, Science, and Society, and is on the advisory board of the Stanford Humanities Center.

- * - * -

OPENING STATEMENT

by

Dr. Arpad Bogsch
Director General
World Intellectual Property Organization
Geneva, Switzerland

Mr. Register of Copyrights, Mr. Ralph Oman,
Mr. President of Stanford University, Dr. Donald Kennedy,
Ladies and Gentlemen,

It is a pleasure for me to greet all the participants in the WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence.

I greet the representatives of our Member States, I greet the observers from various intergovernmental and non-governmental organizations, I greet--with particular thanks for their valuable contributions--our invited speakers, and I greet all the other participants, among them many outstanding experts in this field.

I also thank--through you, Mr. President--Stanford University for hosting the Symposium and--through you, Mr. Register of Copyrights--the Copyright Office of the United States of America for its cooperation in preparing the Symposium.

It is fitting that this Symposium will take place here and now. Fitting, because this area--the famous Silicon Valley--is one of the leading centers of research and development in the field of artificial intelligence; fitting, because Stanford University itself has an important program in respect of both technical and legal aspects of artificial intelligence, and particularly fitting because this is the year when this University celebrates its 100th anniversary.

This meeting is called the WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence.

Let me take the various elements of the title of the meeting, and let me choose a reverse order by starting with the notion of artificial intelligence.

We have defined "artificial intelligence" in our general information document for this meeting as "an expression commonly used to designate those kinds of computer systems that display certain capabilities associated with human intelligence, such as perception, understanding, learning, reasoning and problem-solving."

This is, of course, only a preliminary definition since it is one of the very objectives of the meeting to describe the phenomenon which is called "artificial intelligence," to analyze the existing artificial intelligence systems--such as expert systems (or knowledge-based systems), perception systems, natural-language systems--and to find out in which fields and to what extent it may be considered reality that such systems can perform certain functions of human intelligence, because they are able to perceive, to understand, to learn, to communicate, to reason. From the viewpoint of intellectual property, the most important question is, if such systems can do something that is approaching creation, who should be the owner of the intellectual property in such quasi-creations?

This is why the title indicates that the meeting will examine the various categories of artificial intelligence and their main fields of application from the viewpoint of their possible intellectual property aspects. Discussions will focus, in particular, on what protection should apply to certain existing creations protected by intellectual property rights in respect of their inclusion in artificial intelligence systems (e.g., in expert systems), on what protection, if any, should be granted for the various artificial intelligence systems themselves, and on what the intellectual property status of the outputs of such systems (such as computer-produced creations) should be.

On the basis of the discussions of the above-mentioned intellectual property aspects, an attempt will also be made to assess the possible long-term impact of artificial intelligence on the law of intellectual property (in the light of the fact that existing international systems for the protection of intellectual property have been established with human creations in mind and that, with the advent of artificial intelligence, the possibility of "artificial creation" is not only emerging but, according to some views, already exists in certain forms).

The meeting is called symposium which indicates its nature and objective. The word "symposium" stresses that the objective is to have a free exchange of ideas about this new phenomenon in preparation of further activities in the field that may be found justified.

The adjective "worldwide" is used in the title of the symposium which indicates, at least, two things. First, it indicates that we have invited speakers and participants from all over the world, and, second, it indicates that the subject itself is of a global dimension; the development and utilization of artificial intelligence systems involve close cooperation among researchers, producers and users of all countries.

Hence the interest of the World Intellectual Property Organization in this subject. The acronym WIPO appears in the title of the Symposium. WIPO administers the international treaties on the protection of intellectual property, and it must clarify whether, where such protection is needed in this field, it can be granted under the existing provisions of those treaties, or whether such treaties require modification, or whether some new treaty will need to be concluded eventually.

I am confident that, thanks to what will emerge from these discussions, we shall be in a better position to answer the difficult questions that intellectual property protection poses in the field of artificial intelligence.

- * - * -

WELCOMING REMARKS

by

Ralph Oman
Register of Copyrights
United States of America

Thank you Dr. Bogsch. I and all of my colleagues from the U.S. Copyright Office, the U.S. Patent and Trademark Office, and the U.S. Department of State are greatly honored to participate in this trailblazing symposium.

I heap praise on Dr. Bogsch and his staff at the W.I.P.O.. With great foresight he has organized this meeting to get us thinking about problems and opportunities that await us down the road. I also thank Dr. Kennedy and Paul Goldstein for showering us, not just with the much needed rain, but with kindness and hospitality as well.

Dr. Bogsch has spent his entire professional career fighting a holy war, a jihad on behalf of intellectual property. As a part of the continuing effort, he has assembled this blue-ribbon panel of experts here at Stanford, in the heart of Silicon Valley, to talk about new ways of doing old things. These new technologies will have a great impact on our quality of life in the workplace and on our standard of living.

I'm told that this is the first time he has held a worldwide program outside of W.I.P.O. headquarters. Usually the mountain goes to Mohammed in Geneva. This week, because the matter is so important, Mohammed has come to the mountain.

As always, Dr. Bogsch's timing is perfect. The American pioneer in artificial intelligence, Herbert Simon, has just published his autobiography, and he describes with great excitement the development of these "machines who think." As a result, many people, not just those here in Silicon Valley, but all over the world are talking about artificial intelligence.

For the next three days we will hear about "expert systems," "shells," "knowledge bases," "inference engines," "neural networks," "natural language systems," and "automatic programming."

And you will hear about originality, authorship, ownership, and infringement.

At the outset, let me make just one point. For two hundred years the United States has taken a very pragmatic approach to copyright. We see it as a sort of "bribe," as an economic spur to creativity. By rewarding authors, we encourage them to create. Their creativity, in turn, promotes the progress of science and the arts, and the public benefits. Europe has taken a more idealistic approach. They look upon author's rights as one of the God-given Rights of Man. Books and music spring from the soul of man; they are his children, and they must be revered and protected. With this lofty approach, so different from the hard-headed, practical approach of the United States, our European brothers and sisters in the Berne Union agonize over how to protect works of authorship that might have a useful application--works like computer programs. They may wrestle with the same internal tensions in artifical intelligence.

Over the next three days, I hope we will not wring our hands over platonic absolutes--the fear of diluting the purity of author's rights with an alien concept. I hope instead that we approach it from a non-theological perspective. Let's not ask about moral rights for machines or dwell on non-human authorship.

Instead, let's ask what we have to do to protect the creativity of these "machines who think" in a way that encourages living, breathing men and women to build them and improve them, and that encourages business executives to buy them and invest in their development.

For the next few days we will stand on our tip-toes, peering into the future. The conclusions we reach will help shape that future, and will determine how quickly it arrives.

Thanks again to Dr. Bogsch and Stanford University for making this symposium possible.

- * - * -

WELCOMING REMARKS

by

Donald Kennedy
President
Stanford University
Palo Alto, California
United States of America

It is a great pleasure to welcome all of you to Stanford--especially those who come from relatively far away and expected warm sunny weather and were surprised to see the natives dancing under the shower we had this morning. Let me explain that we cannot worry from our hearts the slightest bit of sympathy for any of you; this is relieving a five year drought and we are delighted with it. So there is one more reason to thank you for being here.

As the Director General has observed, this is our centennial year. We are a young university, in a region in which the matters you will be discussing are very current--mainly because this university has been especially creative and productive in this area. We are fortunate to have such scholars as my colleague Paul Goldstein, who has played an important role in bringing us here, and Terry Winograd--who will appear later on, and is one of several Stanford faculty members who have played important pioneering research roles in this area. Terry and his colleagues are the source of much innovative mischief, and Paul and his colleagues try to sort out the legal consequences of all their innovation. I am sure they can keep him busy for some time!

Intellectual property is the "stuff" of universities. For those of us who try to manage them, Clark Kerr's joke is telling--he observed that universities are only loose collections of faculties united by a common grievance over parking. But to the extent that they are managed at all, questions of intellectual property have much to do with their management. For example, who owns inventions and who owns artistic creations? These matters may have always seemed clear to most people but they have never been treated symmetrically by universities. And now of course, life is getting infinitely more complex. For example, who owns data? This has suddenly become an interesting question, especially if you also consider an additional one: how much do you have to alter a data base in order to claim that you own it? And, of course, the modern science of artificial intelligence, as this gathering will find, is making the matter more confusing still.

Computer scientists have discovered that it is possible to accomplish equivalent hardware and software realizations of the same thing, yet those two are treated under different systems of policy--the patent system and the copyright system. As if that were not enough, it is plain that artificial intelligence and "expert systems" will now provide even more difficult challenges about proprietorship. The Director General indicates that the work of reasoning machines may present us with other new problems, and it leaves one only to wonder how long it will take before the first law review article entitled "Should machines have standing"? is published.

Stanford has always been a place in which the social consequences of technology and innovation have been explored. I remind you in passing that our colleagues on this faculty were the first to propose a moratorium on genetic engineering research--very early in the development that ultimately led to biotechnology--because of uncertainties about its consequences, and that Stanford organized and held the first conference on corporate involvement of university sciences in the biotechnology industry, in the early 1980's.

Stanford is a place that worries not only about what will be next and how we make it happen, but also about the consequences. This symposium is obviously very much in that tradition. So I want to congratulate you for this exploration; to express the hope that it lives up to your expectations, which I know are high, and to voice the wish that I could spend the rest of the day doing what you are doing instead of what I now have to do.

Welcome to Stanford and thank you very much for being here. I look forward to hearing more about what you do.

- * - * -

NOTION AND GENERAL OVERVIEW OF ARTIFICIAL INTELLIGENCE

by

Terry Winograd
Stanford University
Palo Alto, California
United States of America

I am starting off with a task which is, on the one hand, quite enjoyable, and on the other hand, impossibly difficult. I would like to bring you to an understanding that I am afraid many people--including many of us in the field--have not managed to achieve over a very long time--a sense of what artificial intelligence is.

When people ask me what artificial intelligence is, my immediate response is "That depends on why you are asking." It is not a simple question. There is not a simple definition that covers it all, and it takes a lot of thinking and a lot of background to try to situate it at the point where we can not only define it--but also see its consequences, what it can do, what it may do in the near future, and what it might do in the long distant future. That is a very hard quest, which many of us in the field have engaged in, and there are many different opinions. The one thing I can offer you as back-up to the inevitable shortcomings of what I say is a rather extensive bibliography which you will find at the back of the outline I gave out. You will find in those readings a much more substantial discussion of these issues.

Let us first address the direct question--what is artificial intelligence? There are several different starting points that you might take, which lead you to see different parts of it. It is like the old story of the blind men and the elephant. You may see it as something tall and stately or as something long and stringy--it has many different aspects. For most people the phrase artificial intelligence evokes an image of what I call an "eternal quest." This dream that somehow as Homo Faber--the maker of tools--we could go so far as to build a tool which duplicated our own powers--which was able to think, even to create--to be like us in the mental realm just as we have been able to build tools which match or exceed our powers in the physical realm. From that point of view, artificial intelligence is not defined as any particular devices, as any technology, or as any theory, but rather as a goal. What is it we would like to have? I think a lot of the intuitive appeal of artificial intelligence, the thing that makes the researchers want to devote their lives to what is, in fact, rather abstruse technology, is the sense that we are moving towards machines that have that power. I have not yet read Prof. Simon's autobiography, mentioned by the Register of Copyrights, but I suspect that reading that will give a real sense of the feeling of the power of the quest to be able, through the powers of own minds, to produce something which can think.

A different, but related goal is the scientists' goal of understanding how things work. Just as physicists set out to understand the workings of mechanical and electrical things in the world, the people who call themselves cognitive scientists set out to understand what goes on in the brain. Not necessarily to duplicate it, but in the same sense that all science proceeds, to understand it and then to use that understanding to increase our powers to do things. So, a second direction that you can take in defining artificial intelligence is the scientist's view of understanding how intelligence works. In the broad sense it is not specifically animal intelligence, not specifically machine intelligence--but whatever it is that would be common to all those things. This perspective leads to a somewhat different approach, for example, studying in detail how existing thinking systems work. There is no reason to assume that a thinking machine would operate like a brain, any more than a steam-powered sledge hammer operates like my arm. There are some basic principles that are similar, but in detail it is going to be very different. So, someone who is concerned with construction might not care about the particularities of the mental system that we have. On the other hand, we are always fascinated to find out more about ourselves, and there is good reason to suspect that that would be a key in unlocking the secrets of how we can build things.

The third direction one might take is to say that artificial intelligence is defined by an academic discipline which, over the course of the last 40 years or so, has produced a particular set of ideas, of theories, of devices and technologies. It is defined not by the quest, where it might go, but rather by the achievements. To give a mundane example, I am in the process of helping to organize the qualifying exam for our graduate students in artificial intelligence. The qualifying exam does not deal with questions about the quest for creation. It deals with questions like how do you apply second order representations to do non-monotonic reasoning. There is a whole body of technicality, and I do not use that term in a pejorative sense, but rather in the sense of building ideas, constructs, ways of working, methodologies, and devices.

Now, why are all these definitions talking about the same thing? They are not the same because we see tremendous gaps. I want to give you a sense, through this talk, of where those gaps are and a feeling for what they then imply. We are able, using the technology, to build programs which do certain things which, taken in a literal sense, are obviously "intelligent." There are systems that can analyse spectrographs, systems that can perform various kinds of industrial tasks, systems that can answer certain simplified questions in a simplified version of natural language, things which, if you asked somebody fifty years ago, "can a computer do that?"--they would have said, "no, they do computing, not thinking."

On the other hand, we are painfully aware, especially for those of us in the field, of how far removed those tasks are from even the simplest things that we expect people to do. They are so limited and tentative compared to what we expect from a normal six year old, that we cannot say that we could take these dreams of intelligence and these techniques and expect them to come together. Twenty or thirty years ago there was a feeling that we were close to that. I think there may be a few people around who still believe, but the course of time has shown the fact that we can make a step forward in some direction does not imply in a simple way that there is another step and another step and ten more until we come to the place we would dream of being.

Much of what is going in artificial intelligence today is a reflection about where have we gone, what we can do with the things that we know how to build, and where major leaps are still needed--where we do not really understand what is going on. While we can build machines that do incredible mathematical equations, nobody can begin to build one that has the common sense of a dog. In fact there are substantial AI projects now working very hard to build simple robotic models that have the intelligence of a cockroach. That is still well beyond what AI can do. So, we have these marvellous machines doing amazing things, yet cannot reproduce the kind of intelligence that we find in a cockroach. Now, why is that? I want to take a little time on some of the more technical issues at stake.

The mainstream of artificial intelligence work over the last forty years has been in a direction which might be referred to as "symbol manipulation." Professor Simon, who was mentioned earlier, along with his colleague, Allen Newell, put forward something they call the "physical symbol system hypothesis." They argue that the things that we call intelligence are the result of activities going on in the nervous system which are in principle very like symbol manipulation in a digital computer. As we look at our computers and we see that a pattern of magnetic or electronic current represents a set of bits--a 1011--which in turn represents a character, which in turn is part of a representation of a factor, or a sentence, or a piece of data--animal intelligence is very much like that although different in detail. No one expects to find bits of ones and zeros in the very simple sense of the neurons, but at an appropriate level of abstraction, they claim we will find that everything we recognize and admire in intelligent systems that are naturally evolved will be explained in terms quite like explaining how a program works in a computer. There are many, many variants in this hypothesis, and tremendous debates go on as to whether the processing is serial--that is, happening in sequential order--or parallel, many things at once. There are debates about whether the symbol processing--the rules by which the machine takes one set of marks and produces another--correspond to things that a logician would recognize as part of ordinary deductive logic, or whether they have a peculiar logic of their own which, although precise, is not the kind of logic that a mathematician is used to.

There are tremendous debates about the details, but what unites this major mainstream view of artificial intelligence is the hope that we will understand and duplicate intelligence by some device that operates by having well defined rules--algorithms--that operate on symbolic representations. The physical state of the machine is effective because it "represents"--it does not matter that there is a certain set of currents, or a certain set of voltages, or a certain set of magnetic domains. What matters is that there is a standard way in which we interpret that one to be an 'A' and that one to be a 'B' and that one to be a '17', and rules operate at that level. Now, this approach is not something new with computing. For many years, within artificial intelligence, an historical arrogance has existed in the feeling that this is all brand new stuff and therefore it is going to work right away and do everything new that the world has never seen. In fact, when you look back, this view really has been there in various forms all along, probably back to the Greeks, but at least as far as the Enlightenment. There are quotes which, when you read them today, could refer to AI. Leibnitz said that there is no more reason for dispute in reason than for dispute in accounting. When two men have an argument, they should be able to sit down, state their premises, and one of them should say "gentlemen, let us calculate." Hobbes talks about the fact that reasoning is a matter of adding up and substracting parcels just like arithmetic. It is almost a direct quote of symbol manipulation but, of course, centuries before anyone thought that could be done on a machine. I just mention a few names here. The work on logic and the foundation of mathematics in the late nineteenth century and the early twentieth century by computing pioneers like Babbage, established the basic paradigm that says intelligent thinking is a matter of manipulating symbolic parcels according to rules. The computer added to that the ability to actually make it happen on a massive scale. Leibnitz built an adding machine which carried out a symbolic process with physical devices, but it was very crude and could not do very much. Electronics made it possible to do that on a very different scale, so the whole idea which had been an abstraction became a possibility. In the mid 1950's, a tremendous amount of enthusiasm and energy went into trying it out. Now, when you are going out into new territory and you do not know what is there, it is easy to assume that just over the other side is India or whatever it is you are searching for. It must be over that next mountain because nobody has ever been there before and you are headed in the right direction. So, there is an arrogance that we are about to find something new.

What has happened in the last few years is interestingly enough a kind of looking back at alternatives that were left aside in that gold rush, but which took a different approach which now may seem more appealing, i.e., to recognize that intelligence, not only in people but in all kinds of animals down to a very low point on the genetic scale, is embodied in a system of devices--neurons--whose behavior is not at all simple and digital. If we actually look at what goes on in a neuron, there are complex waves and pulses and so on. The notion that there is single firing followed by a "boom" is a gross simplification which has certain uses but ignores a tremendous amount of other structure. They are not homogeneous, and these are all sorts of things about neural systems which are quite different from the computers we build. Notice that the symbol hypothesis says those are uninteresting differences. There are also physical differences between a computer built on a micro-chip today and a computer built out of vacuum tubes back in the 1940's, but that does not matter to the principles of computing. Those differences have to do with reliability, speed and things like that. What it really means to compute--what constitutes valid computation--and what is in a theoretical sense "possible" does not change when you go from vacuum tubes, to transistors to microchips, or for that matter if you were to build a computer out of fluidics and beer cans. What counts is the fact that something represents a certain symbol, an 'a' or 'b' or '17,' not that it represented electrical impulses in vacuum tubes or beer in beer cans, or whatever.

So, the fact that natural intelligence is situated in a body, a physical system which has many other properties, leads to a different kind of speculation which says there is something going on there which needs an explanation very different from the kind of mathematical symbolic logical process that goes into programming AI. It was argued back in the early days that proposed machines were much more like neural machines and much less like symbolic process machines. Given the potentials for the technology, given what could be built, the digital approach was much more promising for us to follow for quite a while. It was much more promising because there was new territory that could be explored within the constraints of what you could afford to build and afford to try.

The last 10 years have been intensely productive of other approaches because two things have happened. One is that the limits of the mainstream approach have become more visible. When I wrote my dissertation in 1970 at MIT, it was natural language understanding--how you get a computer to interact in sentences typed in English and respond to them appropriately. And, it was very clear at the time that what I was producing was greatly simplified. There is no doubt that it was not handling the full complexity of the English language, the full complexity of the kinds of situations that people would talk about, the relevance of context, and so on. The feeling was that simplification was a methodological step; just as if you were studying physics, you would start with something as close as you could get to frictionless objects rolling across flat surfaces. You would not start studying physics by rolling hedgehogs across the lawn; that was simply too messy. So, the feeling was that you get the simple stuff first, but then in a more or less straight forward way you extend outwards.

Now, twenty years later, the state of natural language and understanding systems is not tremendously better. I do not mean to say that it is not some better--there has been a lot of interesting work in many areas related to language--but I think that has gone along with a much more mature understanding of how hard it really is. Context needs to be brought into consideration--the mechanisms that I used, and what the people after me invented, do not do an adequate job of it. There are some special cases but the feeling is very different from what people do and there are no general purpose language understanding programs that perform significantly well. The ones that do work well are very special purpose.

Why is that? One view is that we have not tried hard enough, we have not got the right answer. My colleague John McCarthy here at Stanford, who has won a major prize in science for his work on artificial intelligence, when asked how long it will be till we have artificial intelligence, replies that "it will take something like 2.6 Einsteins and 1.4 Manhattan projects." What he means is this is not normal science, it is not just "fill in the details and you've got it." It is going to take major break throughs that we do not even know yet.

So, given that there has been an interest in looking beyond the central means for areas where there might be a major breakthrough, let me give you a look at the field today. These are just landmarks for going off and reading more. You will see a very interesting mixture of the symbolic mainstream traditions of AI along with things like connectionism and neural networks, in a direction that is called "emergent artifical intelligence." The notion of "emergent" is that you put together pieces without a full analysis of what they will do, and behavior emerges from the whole which is not something that you designed in the individual elements. The behavior of societies is not the same as the psychology of individual people but it emerges from that. Whatever it is that societies do, mobs and armies and political movements cannot have originated anywhere else, but from the people within them. But there is no simple sense in which understanding how an individual works leads you to understand social phenomena.

I want to tie this to the themes of this symposium. First, the most obvious theme is the ability of computers to create. Last week just before this symposium, I got a call from a television station asking me to comment on a story, we are doing a story about a fellow who had announced that he has a computer program into which he has built a model of the author Jacqueline Suzanne--for those of you who are not Americans, she was a writer of popular novels, now deceased--and that his computer had generated a new novel which is just what she would have written next if she had not died. What does it mean to say a computer has written a novel, much less to say it has written a novel based on a model of some other author?

I do not know what that program did. But I do know, from past experience and from general principles, what kind of things it might be doing and what kind of things it probably is not doing. There is no natural language system, for example, that generates coherent natural text. If you look at the output of natural language systems of any kind, unless they are tremendously stylised, it is very rough. They do not take into account subtleties of connotation, subtleties of context. You have to read and puzzle over the output to figure out what it means. Now, what this Jacqueline Suzanne news clipping said is that the man had created it jointly. It had been a sort of joint effort between the person who had done the work and the computer program he had written. It proposed things, then he edited them somewhat. Now there have been examples of this in AI programs, and one of the more famous ones is called EURISKO, done by Douglas Lenat who was here at Stanford at the time, which is not in the natural language area but rather a general problem-solving program. In particular, the problem he tackled was a war game that people play against each other, and his goal working with the machine was to develop new strategies for winning this game. He would input some stuff, the thing would run over-night and come up with many possibilities. The point of the machine was to take the stuff that was there, re-combine it in different ways and elimininate the impossibilities. In the morning, he would come back and throw away the ones that did not look interesting, feed in the ones that did look interesting and let it run again. So we have a kind of mixing of efforts. It is the machine doing a combinatorial kind of problem, that is, taking various possibilities and combining them in ways that nobody had tried combining them before, followed by a real person applying contextual judgement about what makes sense and what does not. In fact, in this particular case, his machine ended up winning the international championship for this war game and being banned from future competition. Partly, because the computer happened to find a loop-hole in the rules. It came across a combination which was so weird that nobody would have tried it. But it turned out that when you did try it, it won all the time, because the rules happened to be structured in a peculiar way.

Notice an interesting property here. Computers are unthinking. They do not have judgement of whether something sounds reasonable or not. That is, the computer apparently has no "sense" of that kind even though it can do the manipulations. The interesting partnership arises when you combine a machine which is so dumb it will put together impossible things with a person who is smart enough to recognize when that impossible thing has some value. So, what I would say in terms of the potential of AI for creativity is that with the techniques we have, I am in the skeptic camp on AI. The book that I wrote with Fernando Flores called "Understanding Computers and Cognition" is a critique of the mainstream AI and philosophical claims that what we think of in the broad sense as human intelligence--judgement, context, understanding--can be replicated through AI. I am a skeptic about what it can do if you take that original goal--to create a tool that thinks like we do. I believe that it may take wholly new systems. One of the early critics of AI, Hubert Dreyfuss, likened it to alchemy. He wrote a famous paper, in fact, called "Artificial Intelligence and Alchemy" in which he said that trying to turn a machine into intelligence was equally foolish as the alchemists trying to turn lead into gold. An interesting sequel to that which I realized long afterwards is that in fact we can turn lead into gold. If we go a few miles from here, there is a two mile long particle accelerator. If you fire the right kind of particles down at a block of lead--you could, in fact, turn some amount of lead into gold. Now, the alchemists therefore, were right. They were right in a very funny sense. When they said lead can be turned into gold, they were not literally speaking a falsehood. But what they were doing was very far removed from what it actually means now, both in terms of their knowledge and in terms of the impact. They would not be interested in this machine that cost a hundred million dollars to turn a few micromilligrams of lead into gold, they wanted it for money.

The same is true with AI. I do not believe it would be grounded to say that we cannot produce intelligent machines in the same sense that somebody in the days of the alchemists would have been grounded and ultimately proved wrong. But we may be as far in our current understanding of intelligence from really building intelligence as the alchemists were from understanding linear accelerators and atomic physics. There is no reason to assume that somehow we are privileged, that it is going to be faster and easier, that this time we got it right.

Now, what does that say for creativity? I do not predict that there will be machines which could be said to go off on their own and produce something creative. What there will be are machines which one person produces, which in conjunction with that person or some other person, do produce things. And where the creativity is and where the production is--where the ownership is--will then be a very funny question, not so much between the machine and the person but between the different people involved. If I write a program and you use it and it produces something great, what is the relation between you and me? AI is a new kind of medium. We must look at the ways in which the programs we produce with artificial intelligence will serve as a new kind of creative medium in which people work, but also one where in some sense the medium is bringing with it part of the game. If I paint a great painting, the person who created the pigments does not have much claim to my creative work. If you stop and think about that, the essence of what was new and creative about this work may have been very much tied up with a new kind of pigment or medium. In some sense it is as much the other person's role as mine. Well, now that we are in the symbolic realm, it is not as clear as the pigment versus the brush stroke, but a much more open question.

So, in conclusion I want to focus your attention on this question of AI as a kind of communication and as a way in which people can collaborate with other people in ways that they have not been able to do so before, across space and time, collaborating with somebody who wrote a program many years ago if I am using it in a way that involves creation.

I wish you all the best of luck in trying to figure out the consequences of that.

ARTIFICIAL INTELLIGENCE AND INTELLECTUAL PROPERTY LAW: SOME PROBLEMATICAL EXAMPLES

by

Professor L. Thorne McCarty
Rutgers University
New Brunswick, New Jersey
United States of America

My plan in this talk is to give you some concrete examples of the applications of artificial intelligence, and then to focus your attention on the intellectual property issues that are raised by these examples.

I will not try to answer the question: "What is artificial intelligence?" But I will state initially a useful characterization of one of the pre-suppositions of AI:

> "Intelligent agents construct internal representations of the external world, and they process these representations in various ways to achieve their goals."

I state this in a very general way, intending to make no commitment as to the nature of these representations--whether they are explicit or implicit, whether they are symbolic, neural, or whatever--and to make no commitment as to the kinds of computational procedures that are involved.

I believe that most people in AI would subscribe to this as a neutral characterization of one of the presuppositions of the field. As to specific questions about the nature of the representations--whether they are explicit, implicit, symbolic, neural, etc.--that is where the schools and approaches and philosophies differ strikingly. What I want to suggest, as I go through some examples of applications that use different kinds of representations and different kinds of computational procedures, is that the intellectual property issues might also differ accordingly.

1. Explicit Representations of Knowledge

I want to look first at some examples of the applications of artificial intelligence that are within the classical AI paradigm, which is based on the use of explicit symbolic representations of knowledge. To date, most of the applications of AI have been within this framework.

To give you some examples, I will first talk about applications that may be of particular interest to this group: legal expert systems. I will give you two examples that are, in some sense, on different points of a spectrum.

(a) Propositional Rules

The first example is represented by a book entitled "Latent Damage Law: The Expert System" by Philip Cappen and Richard Susskind. This may be a forerunner of many books of this sort: You open the back flap and find two floppy disks. The book itself gives you a detailed discussion of the latent damage law expert system--how it works, and what its limitations are. Also, since this is one of the first projects of its kind, the book gives a detailed discussion of how the authors went about constructing their system. The methodology they use is interesting to study. They also have a chapter on some of the legal implications of legal expert systems.

First, what are they trying to do? Latent damage law addresses certain problems that arise when a cause of action--for example, a tort or a contract--is barred after a certain amount of time by the statute of limitations--three years, five years, or whatever. If the plaintiff could not have reasonably discovered the damage or injury when it occurred, there are circumstances in which the statute will not begin to run, and the cause of action will not be barred. The question is: Under what circumstances might you avoid the limitation of these actions?

As I understand it, a new statute covering latent damage law was passed in England, several years ago. It is said to be very badly drafted, and, of course, it must be interpreted against the background of a very complex and confusing body of case law. Apparently, there are serious pitfalls for the ordinary practitioner who tries to navigate through this area of the law.

Philip Capper, one of the authors of the book, is one of the leading English experts on latent damage law. He had previously written a book entitled "The Latent Damage Act 1986: The Impact on the Professions and Construction Industry." It was an ordinary treatise of the usual sort. But then he got together with Richard Susskind, who had previously written a small prototype expert system at Oxford, and they wrote this program.

What does it do? It draws you, the legal practitioner, through a structured interview session to try to determine the answer to a legal question. The question is always something like: "Is the cause of action for X barred at time Y?" It is very easy to describe how this system works and what it does. In fact, it is an example of an expert system that many people in the field would say had nothing to do with artificial intelligence at all. It uses a very simple framework for representing knowledge: propositional rules in the form, "if A and B, then C."

Many of you will recognize this as one of the earliest frameworks for the construction of expert systems. Within that framework, the authors have encoded Philip Capper's extensive knowledge of latent damage law, and the proper way to analyze fact situations involving latent damages. The system moves backwards, from the initial question--"Is my cause of action barred?"--to various subordinate questions. It asks you for answers as it goes, to which you respond by typing "yes" or "no," or by selecting items from a menu. In this way, the system elicits the facts of your particular case.

The system works because Capper and Susskind have very carefully engineered Philip Capper's knowledge of this area of the law into a rule-based format. Philip Capper knows what the explicit legal rules are, of course, and these are embedded in the framework. He knows where the rules are doubtful, and that information is also embedded in the framework; he knows the likely interpretations, and the alternatives, and that information is encoded also. (In such a situation, the user is given help screens and further citations, and told that the problem is particularly difficult.) The system also has a strategy for tracking the user through the case in a reasonable way, without becoming lost. This is actually one of the tricks in writing a useful system of this sort.

Now, all of this knowledge is very different--knowledge about the rules of latent damage law, knowledge about uncertainties and ambiguities, knowledge about search strategies. All of this is encoded in the same simple format "if A and B, then C," with the same simple processing strategy applied to it. Thus, the technical features of this system are not particularly advanced or profound. What is interesting is the sophistication of the legal and strategic knowledge embedded in the rules. From this perspective, the system is a good example of what can be done with very simple techniques at the present time.

To put the point in the terms of intellectual property law: The intellectual product of value here is Philip Capper's knowledge, encoded in these propositional rules.

(b) Knowledge Representation Languages of Legal Domains

Let me now mention a paper of mine entitled "A Language for Legal Discourse."

Many of us in the field of Artificial Intelligence and Law are trying to go beyond propositional rules to develop more sophisticated knowledge representational languages for legal domains. In my work, I am constructing a language in which we can represent some of the "common sense categories" that are needed to describe legal situations and legal rules. For example, we need to represent time, actions, and events, and we need to say whether an action is permitted, forbidden, obligatory, etc. This gives us a basic vocabulary in which we can begin to describe various factual situations, and in which we can write the rules that specify legal consequences. Attached to the specification of the common sense categories are the appropriate inference procedures for each category.

I have previously applied this language to the (U.S.) corporate tax code, and to certain provisions of the Uniform Commercial Code. I am considering applications to some of the more technical rules of patent and copyright law. Another project under development is an estate planning system, designed jointly with Dean Scholbolim, an attorney in San Francisco.

In this project, we are trying to mimic the way an attorney constructs estate plans for various kinds of clients. This seems to depend on "prototypical reasoning." The lawyer initially elicits some information from a client and this suggests a prototypical plan. When the prototypical plan does not quite work, the lawyer modifies it to suit the client's objectives, using knowledge about the rules of the tax code and about the reasons for the plan's failure.

This sort of strategic knowledge about prototypical plans is analogous to Philip Capper's strategic knowledge in the Latent Damage System, and raises the same intellectual property question. But we also have two other valuable intellectual products in our project: There is the "Language for Legal Discourse" itself. This is certainly a major contribution; it is a major intellectual effort to define that language, and get it to work right. But is something that looks like a programming language itself protectible?. There is the specific encoding of tax rules and regulations using the Language for Legal Discourse. This is also a major intellectual effort. Obviously, the tax rules and regulations are themselves in the public domain. But what is their status when they are "transcribed" into a specialized programming language?

I will not try to answer these intellectual property questions at this point. I only present these examples for your consideration later in the conference.

2. Learning in a Symbolic Framework

But let me now be more adventurous, and suggest examples that raise even more difficult questions.

Notice one feature of the examples so far. They are all based on explicit symbolic representations of knowledge, and I have assumed that these symbolic representations are "hand crafted," that is, they are written and encoded explicitly by hand. Now this becomes a very tedious process. One might ask: couldn't we do some of that encoding automatically, that is, couldn't we write a program that would "learn" an appropriate encoding?

To focus our attention, I will give an example in the area of natural language processing: I refer you to a paper by Donald Hindel at Bell Labs, entitled "Acquiring Disambiguation Rules from Text."

If we are doing natural language processing within the symbolic paradigm, we have to write out a grammar of English, say, using a particular grammatical formulation, and we have to write out a lexicon--the definitions of the words--using the terms of this grammar. In the standard approach, we would write the grammar and the lexicon by hand. Hindle has done this in the past. He has written a lexicon of approximately 100,000 words. He has also hand-coded 350 regular grammar rules to disambiguate lexical category. For example, a word such as "sound" in English could either be an adjective, a noun or a verb, and the disambiguation rules are needed to select a particular category in a particular context. But it turns out to be very hard to write these rules by hand. Hindle has worked on this for several years: It is difficult to anticipate the interactions between the rules when you add a new rule, and so on.

The solution is a form of "symbolic learning." There is a text called the "Brown Corpus" that consists of 500 naturally occurring passages in English, including over one million words. The text has been "tagged" by hand, so that every word in the text is classified as to its lexical category. So, the question is: Can we learn the lexical disambiguation rules directly from the tagged text? This is a "symbolic learning" problem in the following sense: We have a symbolic representation language, namely, a way of explicitly encoding lexical disambiguation rules, and our task is to construct a set of rules in this language that will correctly classify the words in the tagged text. Although machine learning is still in a very primitive stage, we know something about the kinds of symbolic learning algorithms that will work on various kinds of problems: One such algorithm was applied by Hindle in his project.

Here are the results. Originally, Hindle used 350 hand coded rules. The error rate on ambiguous words in the text--that is, words that had more than one lexical category according to the lexicon--was approximately 30%. The paper I cited reports several different versions of the learning algorithm, and it is somewhat difficult to understand exactly which figures to compare, but the best I can gather is the following: In the best of several combined approaches, the error rate was 3.5% on the test set and 2.5% on the training set. (In a controlled learning experiment, you take a corpus of 500 passages, you train the learning algorithm on 450 passages, and then you test it on 50 passages. You expect the training set to do better than the test set. The difference between the two is a measure of how much your learning algorithm has generalized.)

So, from an initial error rate of 30% for the hand coded rules, the learned rules have an error rate on the test set for ambiguous words that goes down to 3.5%! Also, the system constructed 12,000 rules, not 350! Obviously, it is constructing very specific rules to pick up certain idiosyncratic patterns of English grammar that appear in the text. It seems very unlikely that you could have constructed these 12,000 rules by hand.

Now, consider our central question. From the point of view of intellectual property law, what is the valuable intellectual product here? Surely, it is 12,000 lexical disambiguation rules. But note: these rules were not written by a person, but by an algorithm (which was written by a person). And further: in some sense, these rules were already implicit in the naturally occurring text in the Brown Corpus.

3. Learning in a Connectionist Framework

My final example is also a learning example, but the paradigm is quite different. There is a fundamental distinction in AI between classical symbolic representations and representations that are based on "connectionist" or "neural network" approaches. One of the arguments for the connectionist approach is that it is fundamentally grounded on a notion of learning.

Let me give you an example of an application of neural networks. This is work done by Mick Noordewier, a colleague of mine at Rutgers, who is a biologist as well as a computer scientist. Noordewier (and his collaborators) are doing DNA sequence analysis. We have a large and growing database of DNA sequences. We know something about the biology of DNA, but not enough. In one experiment, Noordewien is trying to recognize "promoter sites" on DNA. These are sites where the RNA initially binds in the gene transcription process. If you can identify a promoter site, you have identified the beginning of a gene, which is important. Another experiment is to identify what are called "splice junctions." These are points on a DNA sequence in which segments of messenger RNA are spliced out--they are not used for anything--and it is important to know where these occur.

The approach taken by Noordewier and his coworkers is to start with a weak, inaccurate theory of what is a promoter site and what is a splice junction, extracted from the biological literature. They then encode these approximate rules into a neural network, that is, they use the rules to set the network typology and the initial set of weights. I will not attempt to explain in detail the technology behind neural networks, but they basically operate by adjusting weights between "units," some of which are "input units," some of which are "output units," and some of which are intermediate units called "hidden units". (Noordewier is actually using a hybrid symbolic/ neural-net approach, since he starts with a set of approximate rules.)

Given this initial topology and this initial configuration of weights, the network is trained on identified DNA sequences. There are a number of examples of known promoter sites, and there are examples of DNA sequences that are known not to contain any promoter sites. The training phase adjusts the weights, and this refines the theory. The error rate in the promoter site study after learning is 5.3%, and in the splice junction study it is 6.4%. This is considered very good, better than what you could do by taking the best proposed rules in the biological literature and applying them automatically to unidentified DNA sequences.

What Noordewier's system is doing is analogous to learning the lexical disambiguation rules in my previous example, but there is an important difference from the viewpoint of intellectual property law. After the learning, you cannot point to an explicit set of rules that encode knowledge about what is a promoter site and what is a splice junction. The "knowledge" is simply represented by a pattern of weights in the network.

Suppose you think that the valuable intellectual product in this work is the learned knowledge about what constitutes a promoter site and a splice junction. This is valuable information in biology, because you want to be able to perform massive screening of DNA. Maybe you will be able to discover important new genes, and design new drugs--that is why it is valuable. But there is no symbolic representation of this information. You simply have a pattern of weights.

I suggest that the difficult question is: should intellectual property protection apply? And if so, to what, exactly, and how?

4. <u>Conclusion</u>

I have given several examples of AI applications in this talk, with very different characteristics. These include: (1) explicit representations of knowledge, in two forms; (2) symbolic learning; and (3) neural net learning. I have tried to show that these are very different paradigms within the overall field of artificial intelligence.

Further, I have tried to suggest that these differences in paradigm will make quite a difference in our analysis of the issues arising under intellectual property law.

References

1. Philip Capper and Richard Susskind, Latent Damage Law: The Expert System (Butterworths, 1988).

2. L. Thorne McCarty, "A Language for Legal Discourse. I. Basic Features," in Proceedings, Second International Conference on Artificial Intelligence and Law, pages 180-189 (ACM Press, 1989).

3. Dean A. Schlobohm and L. Thorne McCarty. "EPS-II: Estate Planning with Prototypes," in Proceedings, Second International Conference on Artificial Intelligence and Law, pages 1-10 (ACM Press, 1989).

4. Donald Hindle, "Acquiring Disambiguation Rules from Text," in Proceedings, 27th Annual Meeting of the Association of Computational Linguistics, pages 118-125 (1989).

5. Michiel O. Noordewier, Geoffrey G. Towell, and Jude W. Shavlik, "Learning to Recognize Promoters in DNA Sequences," in Proceedings, 1990 AAAI Spring Symposium on Artificial Intelligence and Molecular Biology (1990).

6. Michiel O. Noordewier, Geoffrey G. Towell, and Jude W. Shavlik, "Training Knowledge-Based Neural Networks to Recognize Genes in DNA Sequences," in Advances in Neural Information Processing System, Volume 3 (1991).

THE MACHINE TRANSLATION SYSTEM AND ITS POSSIBLE INTELLECTUAL PROPERTY IMPLICATIONS

by

Makoto Nagao
Kyoto University
Kyoto, Japan

There are a number of unique computer programs in the field of artificial intelligence. Special features common to those programs are that they are designed to do complex human work, so as to replace or help human intelligent activities, and to produce results widely valuable for society. They are usually designed to have a learning ability in which systems can learn and gradually improve by themselves. However, in some cases, these systems produce wrong output for a certain kind of input data, and in most cases, they need close interaction with humans, or human intervention, to perform an intelligible job. In this way, the nature of artificial intelligence systems is completely different from that of ordinary software or hardware systems.

Ordinary software has been discussed in detail in the context of copyright and intellectual property rights. In addition, a certain framework has already been established for the treatment of software. However, software in the area of artificial intelligence is so different from ordinary software that we have to discuss the problem separately. Therefore, I would now like to focus the following discussion on a branch of artificial intelligence systems, or namely, the machine translation systems that I have studied for many years, in connection with the problem of intellectual property rights.

1. The Structure of Machine Translation Systems

A machine translation system is a software system designed to translate sentences (Sa) in one language (a) into equivalent sentences (Sb) in another language (b). The system usually consists of a program which executes the translation process and linguistic information on languages (a) and (b). The linguistic information includes a dictionary and a grammar used to analyze sentences in language (a), a transfer grammar used to convert the result of structural analysis in language (a) into a structure in language (b), a bilingual dictionary used to transfer words (including idioms and phrases) from language (a) to language (b), a grammar used to generate sentences (Sb) from the internal structure in language (b), and a dictionary of language (b). Grammar rules and lexical information included in this information are interrelated in such a way that as one becomes more complex, the other becomes simpler. Some systems have a large volume of lexical information but a very small number of grammar rules, while other systems carry hundreds to a thousand or more grammar rules for a single language.

2. Development of Machine Translation Systems

It takes a long time and a large number of people to develop a machine translation system. The program for the translation process is usually designed to be able to handle many languages, but it needs to be improved continuously after development is completed. This also is true in the case of ordinary software. Grammars and dictionaries are prepared by linguists, grammarians and translators who specialize in the languages to which the system will apply. This work also takes a considerable nunmber of years to be finished. Meanwhile, languages have countless exceptional expressions, and new words are created continuously. Therefore, to cope with increasing varieties of expressions, the grammars and dictionaries also need to be constantly improved after the system is in operation. After the system is sold to customers, these improvements are made regularly as maintenance work to the customers.

3. Operation of Machine Translation Systems

Present machine translation systems are not capable of producing output of sufficient quality. The quality is certainly improving, but it is difficult to speculate whether any of the systems in the near or even distant future will advance to the point of performing translation of sufficient quality for any kind of input texts. However, it will be possible to develop a system capable of constantly translating some restricted type of texts at a certain quality level, because the quality of output depends not only on the performance of a system but also on the quality of input texts. For these reasons, the operation of the present systems generally includes the process of human pre-editing (rewording or partially correcting input texts) or post-editing or post-revising (rewriting or partially correcting output). Even though such human assistance is required, it has been reported that there are cases indicating that computerized translation is more efficient in terms of time and cost than human translation, especially when technical documents are processed. We expect this style of operation (system use with human assistance) will be more widely employed in the future.

4. Intellectual Work Involved in the Development and Operation of Machine Translation Systems

(a) There is no established or standard algorithm for the translation process at present. This is because there are no established theories or specifications on the design of grammars and dictionaries and, therefore, we cannot determine a universal algorithm for utilizing the linguistic information. Every system has its own characteristics in design. These characteristics are so unique that it is difficult to think they will be integrated into an established algorithm which can be a type of standard.

(b) The situation is the same for grammars used in translation. There is no established grammar for any language. Languages have no such universal or stable laws as physical or chemical phenomena in natural science. Grammar rules vary with linguistic theories. A detailed grammatical explanation will be possible for some limited phenomena observed in a language, and some grammatical rules will be given for these linguisitc phenomena. However, no grammars have ever been written to cover the entire language in the field of linguistics.

In such a situation as this, each system has its own grammars written according to its design by the people hired for the construction of a system. They are usually not specialists in linguistics but have certain training in grammar writing. The grammars written by them must be improved constantly to cover linguistic phenomena as widely as possible. This situation will not change in the foreseeable future. Grammar writing is a trial and error process.

(c) Similarly, dictionaries are prepared according to the design of individual systems. In the Japanese Electronic Dictionary Research Institute, however, neutral dictionaires to be applied to any type of system are being developed under a government-aided project. These dictionaries are designed to be used by any machine translation system. However, if additional information is necessary, it must be added to the original information by the people responsible for constructing the system.

(d) Expressions in the output sentences from a machine translation system are determined by the information included in grammars and dictionaries. The word order of a sentence is determined mainly by grammar rules, while translation words are selected based on information in the dictionaries. As word dictionaires are improved to include more and more units of grammatical information, the quality of translation becomes more dependent on the quality of lexical information, or rather, what is included in the dictionaries. When certain words in a sentence are translated very well, it is often the result of elaborations in entries corresponding to translation words (phrases) in bilingual and target language dictionaries. However, this it not necessarily true in every case because the quality of translation also depends on the quality of grammars used to analyze input texts and generate output. In some cases, it is difficult to determine whether grammars or dictionaries are more responsible for the quality of output.

(e) Dictionaries supplied by machine translation systems manufacturers are those prepared for commonly used words. Dictionaries of technical terminology in specific subject areas often must be purchased separately from the manufacturers of the systems in use or from other suppliers. There are often many special terms which appear in specific texts of a specific user. In such cases, the user must feed translation words in a target language into their computers. For these reasons, machine translation manufacturers usually categorize dictionaires of their systems into general word dictionaires, terminology dictionaries and user dictionaries, then impose restrictions on users on which dictionaires or part of a dictionary can be revised by the user.

(f) As was mentioned earlier, the quality of output is not at a sufficient level in the current systems. In many cases, manufacturers supply additional software for pre- or post-editing. People who use these functions are called pre-editors and post-editors (or post-revisers). Post-editors correct output as much as necessary. Required revision varies from very minor correction to complete rewriting. Post-editors are usually considered responsible for the whole translation text, even if there are many sentences which are not actually revised. Therefore, when output is provided to readers after post-editing, it is not necessary to specify what system has been used. However, if no editing is performed, it is appropriate to provide output to the final reader with information on a system being used.

5. Prospects for Machine Translation Systems

(a) Although the present translation systems do not have a learning ability, they will probably have an automatic learning function similar to word processors in the near future. When more advances are made in research, systems will be able to revise not only dictionaries but also grammar rules by automatically learning from results of post-editing, and they will also be able to produce output of high quality for various styles of texts. When the quality of translation is enhanced in this way, we will need to evaluate properly the quality and quantity of work users have done on the systems in the process of automatic learning. The intellectual property right will be very complex at the time when the dictionaries or grammars of such systems or the systems themselves are provided to other users.

(b) The eternal objective of basic research on machine translation is to develop a system capable of performing a translation task in the same process as human translation, or in other words, "understanding" a text before translating it. The system will be designed to have general knowledge and common sense in addition to linguistic information so that even metaphorical expressions can be translated correctly according to context. The system will naturally be equipped with an automatic learning function, and therefore how this function is used and how the system is "trained" after the installation to a user site will be a major factor for the evaluation of the system's translation ability. We will have to wait at least until the early twenty-first century to see such a system realized. Having an automatic learning function does not mean that the system needs no human assistance. It will be necessary to instruct the system on what translation is right or wrong and to continue feeding examples of translation for automatic training. Such examples will be selected according to the subject areas and styles of texts the system is expected to produce. If the system is improved so extensively that it can automatically select these training data, it will certainly need to be allowed to have "human rights" as an individual, though no one knows when such a situation could be realized.

6. Connection with Intellectual Property Rights

To discuss the problem of intellectual property rights in machine translation systems, I would like to analyze this from the following points of view:

- Grammatical and lexical knowledge in the field of linguistics, and knowledge concerning software,

- Programs and data, which are the engineering reformulation of the above knowledge in a machine translation system,

- Grammatical and lexical knowledge acquired automatically by the system itself,

- Rights to translation texts produced by the system.

(a) Grammatical and lexical knowledge in linguistics is described in books on linguistics or grammar, and dictionaries. Theories on the design of machine translation systems are also presented in books or scientific papers.

This information is all theoretical in explanation, similar to that of a solving method for a mathematic equation. Therefore, as they are, they cannot be used directly in the designing work of a practical system. People engaged in the development according to the design framework of the system must write grammar rules and lexical data using their own knowledge. Dictionaries and books on linguistics and languages are only references for their work. Therefore, we can say the system developed in this way is an intellectual product of all the people involved in the development work, and the reference books and dictionaries cannot claim their own intellectual property rights for the system.

(b) When a large volume of texts are given for machine translation, a system makes various errors and even fails to produce sentences. Post-editors or people in charge of the system management check these errors and feed appropriate translation words or sentences into the system. As a result, the system is gradually improved. Basically, persons who have developed such a self-improving mechanism for a system, and those who assist the system in executing the learning process by using a variety of exemplary texts, are considered as participating in the improvement of a system. To go further in this learning procedure, we can think of a completely autonomous learning process. That is, by the assumption that a large volume of bilingual texts are prepared for learning purposes, the system itself selects a proper set of texts among them to improve the ability of its own system. In this case, most of the intellectual property rights to the system will be shared among people who designed such a self-improving mechanism and who prepared the bilingual texts. In order to improve a system efficiently by automatic learning, the system must be equipped with such a function to select intentionally from the available bilingual text data only the appropriate exemplary texts for the system's purpose.

(c) Post-editors are responsible for the output text they revise. When output text is provided to readers without having been post-edited, the manufacturer of a system in use may have responsibilities and rights to it. The translated text may sometimes cause trouble for the readers. An accident can happen when someone relies on a translated text, and acts as it is said in the text. In such a case, a serious problem will arise concerning who should take responsibility for the accident or damages. If critical errors happen as the result of automatic learning, it will be more difficult to judge who is responsible. I wonder if we will have to handle such a case in the same way as parents take responsibility for their children's actions.

(d) If responsibility could be taken in that way, it may also be appropriate to think that an artificial intelligence system could be allowed to have legal rights and obligations similar to those of humans, when it is regarded as "matured" through experiences in the automatic learning process. Once we approve the "human right" to such a system, we will be able to solve complex legal problems expected to arise. In that case, we have to allow systems to have their own "income" for services they render, and to pool money for insurance in case they cause accidents and damage. In the future, we may even face a problem of the legal right of inheritance to a system itself, or to the knowledge owned by a system and a mechanism to use it.

7. Conclusion

Many of the characteristics of machine translation systems I have discussed will be applicable to systems of automatic programming. Computer graphics programs and automatic music synthesizer programs have somewhat the same nature. When expert knowledge is provided for a certain expert system by a specific person, it will be appropriate to consider that rights to and responsibilities of the system should be shared by the person who provides the knowledge and the person who represents that knowledge in a form suitable for the system, such as production rules.

We can imagine an ideal robot easily which is perfectly autonomous. However, at least for the next fifty years, for the standpoint of legal rights, artificial intelligence systems will be in the same state as is described in this paper. Researchers in artificial intelligence aim at such an ideal goal, but the legal framework for the products and production activities must be based on reality.

MAIN CATEGORIES OF ARTIFICIAL INTELLIGENCE AND THEIR INTELLECTUAL PROPERTY ASPECTS

by

Andy Johnson-Laird
Johnson-Laird, Inc.
Portland, Oregon
United States of America

When I first started doing the research for this presentation, I had a rather nasty shock. I picked up the January/February issue of "BYTE" magazine, and on the front cover was the title "Is AI Dead?" I thought if it is, this is going to be the shortest seminar on record. As it turned out, from the point of view of intellectual property law, the converse of that question is probably more interesting, namely, is artificial intelligence alive?

Let me back up a block, as they say in programming, and rather than thinking about artificial intelligence straight away, I started to wonder how you would measure real intelligence. So I asked some cognitive scientists that I know--what is real intelligence? And they said we don't know. Some of the offerings that I had were, first that "it is the ability to learn and to infer;" second, more intriguingly, that "real intelligence is the ability to explain why you are wrong," and, apparently, "to avoid doing it a second time around."

For the attorneys in the room, I presumed that you would prefer to have some simple, concise tests to use for artificial intelligence, not the least of which is one that is known as "Turing's Test." Basically, this tests whether you can trick a human being into believing that a computer is actually intelligent. If you put a person in one room on a terminal and a computer in another room, can that computer respond in a way that the person at the terminal would say "this an intelligent being." Of course, that begs the question--how do you select the human in the test? If that human is not particularly intelligent, he will be easily fooled.

There is a hundred-thousand dollar reason why this should be a valid test, because it was reported in last Tuesday's Wall Street Journal, that a Mr. Lobar had put up a prize for the first person to create a program that could pass the "Turing's Test." Personally, I also like the definition that says "artificial intelligence is not real. Machines do not think, you only think they think," and there is a corollary of this claim which is "... we think." You can also use the Duck Test--if it appears to be intelligent, it rubs shoulders with things that are intelligent, it walks in an intelligent way, it is intelligent--that is artificial intelligence.

The real point is that we start this conference without an adequate definition of intelligence--real or artificial. As for the question that was asked earlier regarding the difference between AI and software, I do not think, based on my research, that people really know the answer. This problem exists because the world of computers, including software, is a giant swamp.

It is very soft, it is very squishy, and it is very hard to take a firm stand on some of these issues. In fact, as far as I can tell, intellectual property law only exists in those areas (like islands in the swamp) that are sufficiently firm to support the weight of intellectual property law. And artificial intelligence is merely one of those newly emerging islands in the swamp. People are starting to take notice and say, "this is important--how is it going to impact on intellectual property law?" Artificial life is yet another emerging island in this swamp, and is probably worthy of consideration for another symposium.

Let me give you a broad overview of the major categories of artificial intelligence as I see them. We have already heard talk of expert systems. Genetic algorithms were mentioned briefly. These are programs that learn by deliberately mutating. You could also argue that computer viruses are a form of artificial intelligence. They certainly know how to replicate and travel. We have only seen what are marginally malignant viruses thus far. You can also imagine a class of viruses that are benign--you could create a virus that you inject into the world's computer networks to go and find a person's phone number. This virus would self-propagate through the networks and when one of its kind found the phone number it would come back to you and say "here it is." All those outgoing copies that were unsuccessful would simply die a graceful death out there.

I think from the intellectual property point of view, as was discussed this morning, we have two major classes. The first is symbolic knowledge representation, where there is a kind of one-to-one correspondence between real world concepts and internal data structures, and you can look in the internal data structures and say, "yes, what that represents in the real world is this." An example is rule-based systems, where you can look at the individual rules, even after they have been compiled or translated, and you can determine that from those data structures there is a one-to-one, or one-to-many, or many-to-one, correspondence. You can at least look at the internal data structures and understand what is going on.

The other class is non-symbolic knowledge representation, such as neural networks, and other cognitive systems. Not even the author of a neural network can tell you, once the network has been created, what the data structures actually mean.

There is another way to look at it from the intellectual property point of view, and again there are two classes of AI systems. The first is the so-called regurgitative systems--those where you feed in knowledge and rules into the system, and then ask the system questions. The system applies those rules to the knowledge, out come the answers. But the essential point is that it cannot answer questions if you have not already told it the answer. The knowledge is fed in, the knowledge comes back out. Smart information in, smart information out. Also dumb information in, dumb information out. As far as I can tell, those kinds of systems appear to offer relatively few challenges to intellectual property law. They seem to be covered by the existing branches of copyright, trademark, and patent.

Far more difficult questions characterize the other class, the generative systems. You train these by presenting facts; they learn from you or the data that you present them, correcting the errors they make, and they can generate information that you did not tell them in the first place.

I propose to give you another view of one such class of generative systems, the so-called neural networks.

Fig. 1

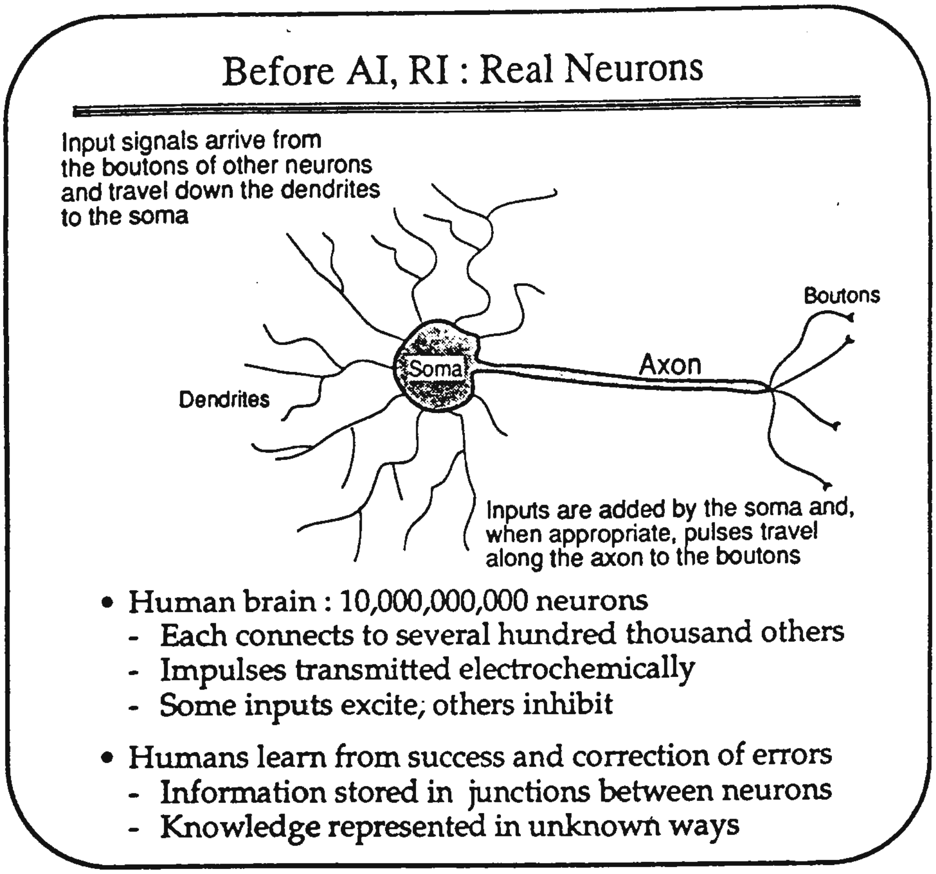

Let me give you a brief introduction by talking about real intelligence and real neurons (See Figure 1). This thing that looks like a surprised squid is a brain cell, a neuron. You have, one hopes, at least ten billion of these between your ears right now. In human beings they are very small--fractions of an inch. In a whale they may be thirty feet long.

What happens is that electro-chemical signals are received by these neurons. The soma serves to add them all together and figure out what to do. If the circumstances are right, a series of electrochemical pulses are transmitted along the axon to the so-called boutons and thence to dendrites of other neurons. Some of these dendrites serve to inhibit the activity of the brain cell, some serve to stimulate it. As I said, you have ten billion of these. Each neuron that you have connects to about a hundred thousand others. It forms a giant network. Humans use these, we hope, to learn. In general, human beings can learn from both success and failure. Neural networks are not so sure; they only seem to learn from failure. When they are wrong you tell them the correct answer and they adjust. When they are right, it is not clear that they are actually learning.

Fig. 2

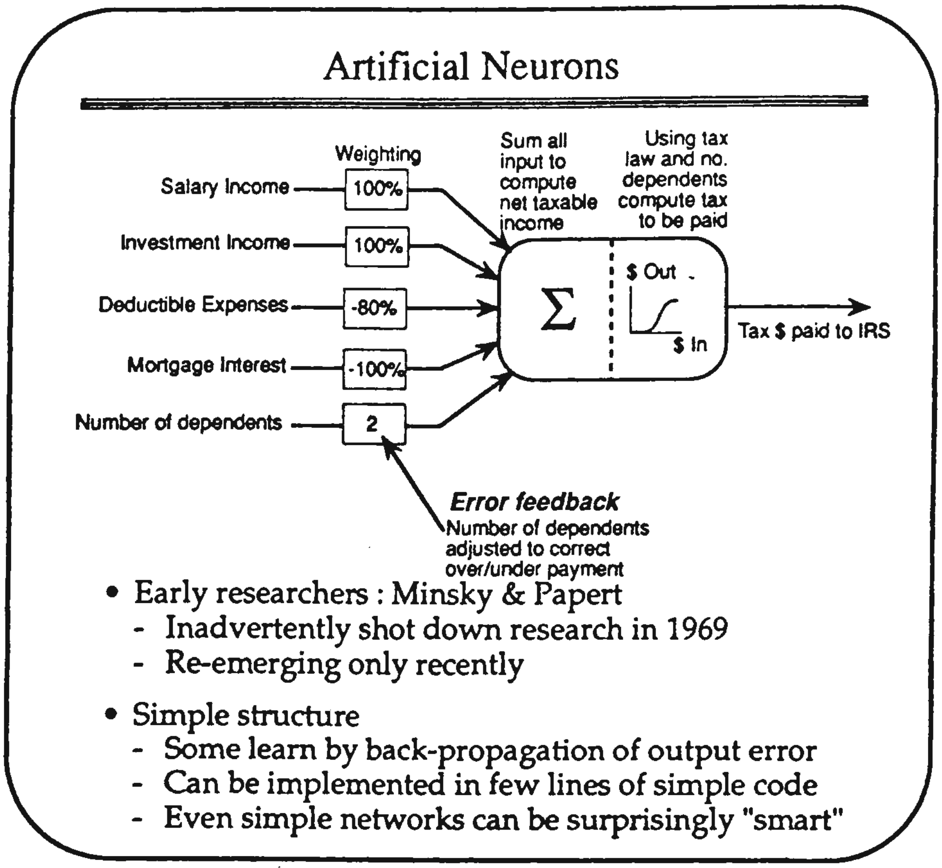

I will show you an example of a software neuron, an artificial neuron (See Figure 2). I am going to use what might be viewed as an odd example because it talks about the U.S. tax structure, served by an organization which is not noted for being endowed with neurons. The structure is repeated. This corresponds to the soma. Here are those factors that affect the amount of tax that you are going to pay. They have variable weightings. Sometimes you will have a hundred percent of your salary income affecting the amount of tax that you are going to pay. Mortgage interest, at least in this country, is deductible under certain circumstances. That tends to inhibit the amount of tax you pay. With investment income, one hundred percent passes through. The number of dependents that you declare will dictate, again, how much tax you pay. This is all summed together and a transfer function, as it is called, is used to determine on tax tables your tax bill, and out comes the amount that you pay to the IRS. This number of dependents is interesting--if you pay the wrong amount of tax, that error is fed back and used to adjust the number of dependents--it essentially becomes a method of regulating the amount of tax. This is an example of a system that learns. If you get the answer wrong, you adjust one of these so-called weights, until the right answer comes out--the system is learning.

AI is not some new branch of software, it has been around for a long time. This particular branch has enjoyed a resurgence. Back in 1969, a couple of very eminent researchers (Minsky and Papert) said "forget it chaps, this line of research won't fly," and they inadvertently shut down funding.

If you read a copy of their book called Perceptrons, in the second version there is a very interesting disclaimer where they said in effect, "sorry we didn't mean it." They have changed their position; they have learned some new information, and now neural networks are becoming extremely popular. In 1991, just about every major magazine in the computer industry has had articles on neural networks, and therefore it will not be long before people start to fall into disputes over products that use neural networks.

Neural networks themselves have extremely simple structures, and even networks that may only require two or three hundred lines of software code can be remarkably smart.

Fig. 3

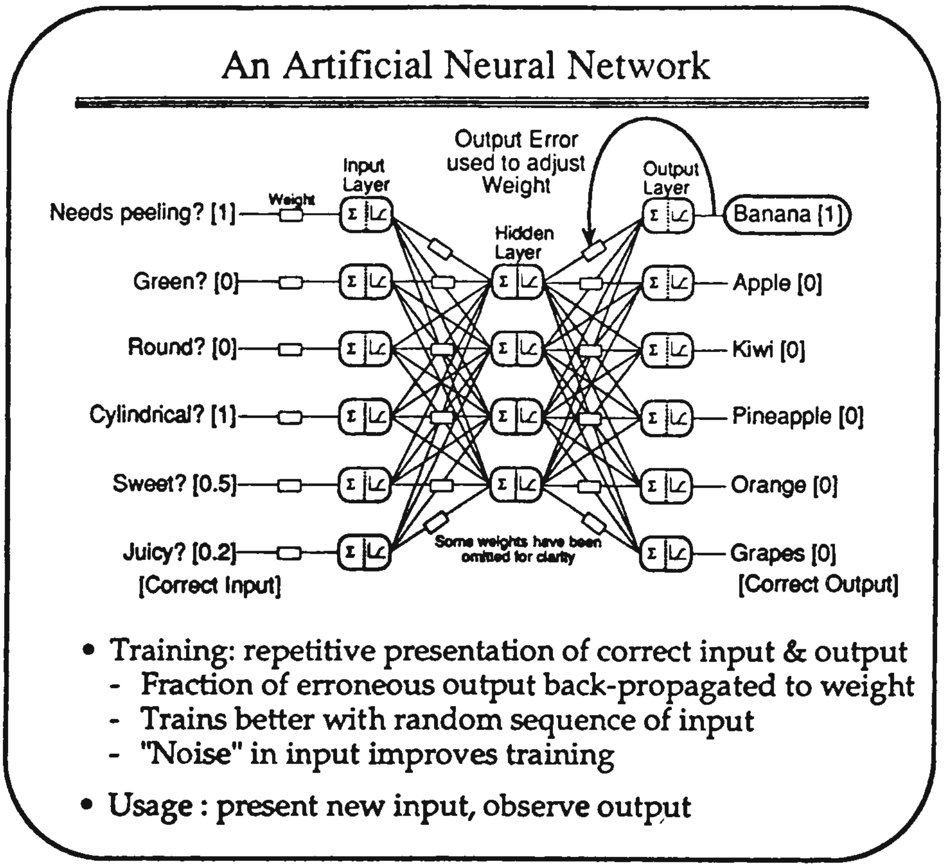

Let me show you one such neural network (See Fig. 3). This happens to be an example that would recognize different kinds of fruit based on certain parameters that you tell it. Here you will see the various input neurons, each with its own weight to adjust the effect of a particular input parameter. Next is a hidden layer of neurons. Each neuron in this hidden layer is connected to every neuron in the input layer. It is this hidden layer that was the problem in the early neural networks. In the early seventies, without this hidden layer, there were certain restrictions that effectively caused Minsky and Papert to shut down research. Here is an output layer. Each one of these neurons represents a similar kind of structure as before. They sum their input, they figure out what the transfer function is, and out comes the output.

How do you train these neural networks? You present them with a series of facts. One fact would say, for example, this particular fruit needs peeling. I am going to represent that as a number one. Is it green? No; so zero, and so on. You can build a series of facts, you can encode them, in any way you choose. The neural network does not care. You then present the neural-network-in-training with the correct answer. So you say, here are the facts, this is the answer. And each mistake, each error is then back-propagated through the network to adjust these particular weights until the right answer comes out, just like the number of dependents in our earlier example.

So by presenting a whole series of facts and presenting the right answer, the weights gradually adjust, until you can take the entire set of the training data, feed it to the neural network, and it will get every answer correct. A couple of curious things happen. Firstly, if you present the facts in an organized fashion--certainly with this class of network, called a "back propagation network"--the network takes an inordinate amount of time to learn. If you take the facts and shuffle them, the network will learn better and faster. Also, if you introduce a certain amount of noise, that is, errors, it learns faster.

To train a network you must run through these hundreds of thousands of training runs--which may take, depending on the speed of your computer, days or even weeks. Once you have a trained network, you can use it by presenting it with some new input that both you and it have never seen before, and merely by simulating those neurons, out comes the answer. It is at this point that you will notice the curious phenomena that the network can now generalize. If you trained it with all the parameters that described a banana correctly and now you tell it about a green banana it will still say it's a banana, even though there is an anomaly in the input.

One of the other curious characteristics of neural networks is that you can take exactly the same structure of network--we call it topology (to avoid being too obvious)--for a completely different application (See Figure 4). This figure shows a network for processing loan applications with the same topology as my previous example. As input parameters you would feed in income, whether or not you are a home owner, occupation code, zip code, whether or not you have a law degree, etc. This same network can then be trained to analyze the input parameters and determine what kind of loan risk you run. As an interesting aside, what I just described to you is illegal in this country. If you use the zip code as a basis for granting or withholding a loan, you are "red lining." And yet, if you also read the January [1991] issue of PCAI, in it there is a neural network described that uses the zip code. So, arguably it is not just intellectual property law that will be affected by this kind of software.

The key point is that the same structure of a neural network can be used for a completely different application than the one that you put into the network. How you encode it is up to you, and how you get the answers out is also up to you. The codification is clearly in the eye of the beholder.

Fig. 4

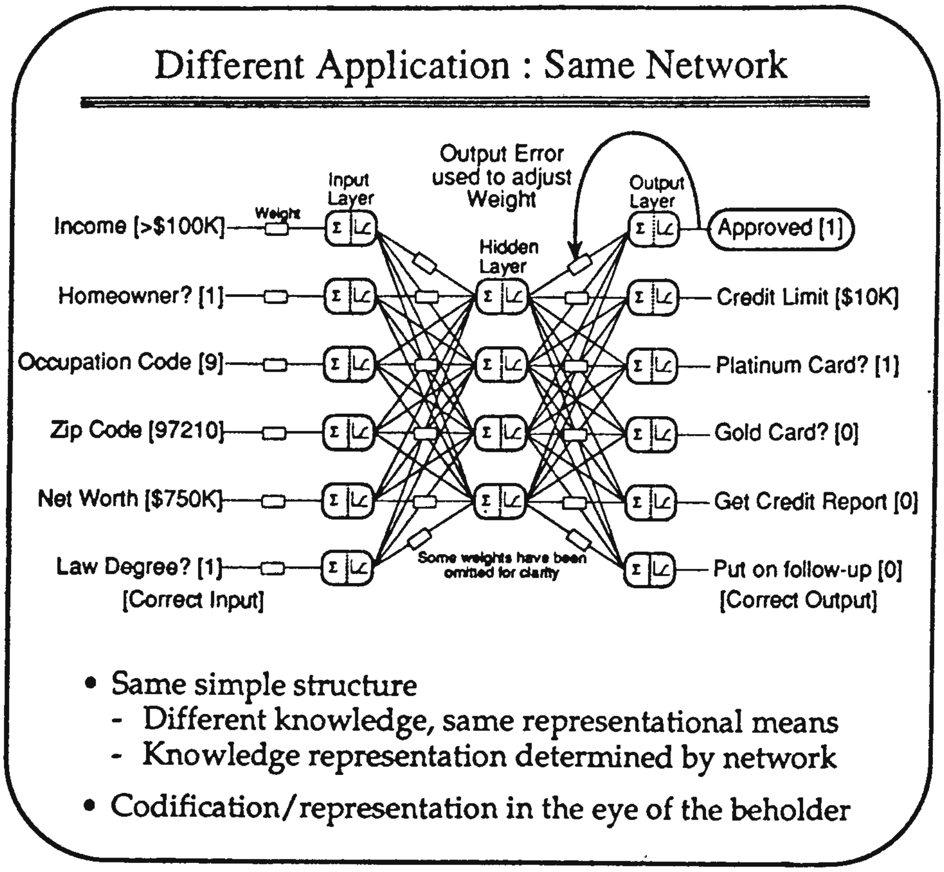

The legitimate question is: where is the knowledge held in all of these systems? Where does this actual knowledge reside? And the answer is: it is in these mathematical weights between the neurons. They correspond in human physiological terms to synaptic junctions between one neuron and the next, which is, I suspect, where knowledge is stored in human brains.

When you start with a blank neural network, you initialize those neural weights to pseudo-random numbers. (The network trains faster with pseudo-random numbers.) Yet even when you have trained the neural network with hundreds and thousands of facts, you end up with a group of pseudo-random numbers.

If you took the same training facts and shuffled them yet again, and re-trained your network on the same facts but in a different sequence, you would get a different set of pseudo-random numbers. In a sense the information is held in a distributive form throughout the entire neural network. You can actually go in and deliberately damage the network, change one or two of the weights and, rather like a human being who has suffered brain damage from a stroke, the network will get by. It will learn, it will re-adjust to that partial damage. It is rather like a holograph. If you hold a holograph up to ordinary light, you see a uniform grey pattern. The entire data in the holograph is stored throughout the holograph. You can break the holograph in two and illuminate it with laser light and you will still see the same image.

In general, we cannot look at the neural weights and discern what knowledge is stored in them. The best we can hope to do, at the moment, is to waggle one neural input and see what it does to the output. You can say, "that is very sensitive, maybe the zip code is important in terms of whether or not you give a loan application."

Here are the questions and challenges I see for intellectual property law, copyright, and trade secrecy.

First, authorship--who authored the knowledge in the neural network, the neural network itself, the actual code that simulates the behavior of the neural network? This is a question that is easy to answer. It is the person who actually wrote the code that simulates the neural network.

But what about facts used to train the neural network? If they are facts, can they be protected? Well you can say, I have organized them in a particular way, I would like to protect them as a compilation. But the networks learned better when you randomized the facts. How could you protect those? I talk here about the ownership of knowledge, and I know, under the law, ownership and authorship are fairly close together. What I am talking about now is someone who produces a trained neural network, puts it onto the market and allows further training by the user. If you start with a hundred lines of source code licensed from somebody else and you add to it and end up with two hundred lines of source code, then you clearly have a legal interest in that software. However in the case of the neural networks, if you add a great body of knowledge to a network that has four thousand neural weights in it, at the end of all that training, you still have four thousand neural weights, albeit of different values. However, the way they have changed value is by going from one pseudo-random numbered group to another pseudo-random numbered group.

In terms of infringement, it is going to be exceedingly difficult for an expert to look at the neural weights and tell you what they represent and whether there is plagiarism. If anything corresponds to object code or machine code (in terms of the classical terms of software--source code and object code), the neural weights are probably the closest we have. Mr. Ficsor asked the question "can it be downloaded, can it be dumped?" And the answer is "yes," just as easily as object code. You can also have one neural network and connect up another neural network to it and siphon off the information by randomly setting the input neurons of the original network to all possible combinations and observing what comes out, and training your own neural network to give the right answers based on this. You cannot actually decompile the neural weights because they are pseudo-random numbers. You do not know what the representation is. You cannot just disassemble them, nor can you necessarily find any meaningful patterns in them at all. In fact, if you did find patterns in them, the odds are it is a by-product of mistraining.

That raises the question, how do you look for possible infringement? If I have a neural network that specializes in loan applications, and I spent literally years training it with a huge body of knowledge, and I see someone come onto the market with a similar neural network, how can I detect infringement? Could I use the cartographic trick of putting in false roads and false lakes in the form of false knowledge? Could I train my network to answer such questions as "What do you think of the Prime Minister whose parents are trapeze artists?" If an allegedly infringing network knows the answer to this question, could it be argued successfully in court that this is a clear indication that the knowledge in the original network was stolen? If I put in four thousand rather strange facts and I can show to a judge that there are three thousand rather strange facts that all match, can I argue that infringement has occurred?

You can certainly apply the conventional test of access to the original work, and you can apply the criteria of reasonable effort, but you must keep a meticulous audit trail of your efforts to be able to demonstrate that these are your training facts, and here are the logs of all of the training runs that you have actually executed, in order to be able to substantiate that all of that information came from the sweat of your own brow.

You must also consider questions of the ownership of neural network output. Does the current view of machine generated work still apply? How will a poetry be viewed, when it is emitted by neural networks or generative artificial intelligence systems? What about wisdom or knowledge, in the sense of advice? Who would be liable when a bridge falls down, if it was built on advice from a neural network, or a generative piece of AI software?

I have not really said too much about software patents in this context but that is another swamp, another minefield. Certainly from my view as a software practitioner, we are in a terrible situation. A question of disclosure was raised earlier. How can I possibly generate a piece of software now, only to discover three years down the road that I have infringed? It seems to me there is an urgent need to view patent law, copyright law, and trade secrecy in the context of these generative works.

What we do, or what, I should say, you do, as members of the legal profession, remains to be seen. Maybe we should consider sui generis protection for this class of software. It is so strange, so different from what we have seen before. I am not sure that existing law could be stretched to cover it. Certainly, we have got to try and avoid adding some of the mythology, that has already sprung up around existing software development. We have to prevent that legal mythology springing up around this class of software. As far as I can see, we are on the right path at least. This swamp must be drained. We have to get rid of this sea of ignorance. When you drain it, you will see that all the islands are connected, and that there is a continuum of software ranging from AI to micro-code. Even though it is not easy to drain that particular swamp, to coin a phrase, perhaps, the toughest part is to drain that swamp when you are up to your backside in litigators.

COPYRIGHT PROTECTION FOR ARTIFICIAL INTELLIGENCE SYSTEMS

by

Morton David Goldberg
and
David O. Carson
Schwab, Goldberg, Price and Dannay
New York, N.Y.
United States of America

I. Introduction

Among the intellectual property aspects of artificial intelligence ("AI") to be considered at this conference is "what protection, if any, should be granted to the various artificial intelligence systems themselves."[1] The statement of the issue suggests that there is some question whether AI systems are in relevant respects different from the works that previously have been recognized as forms of intellectual property.

It would be presumptuous for lawyers to pretend to resolve whether AI is qualitatively different from previous works,[2] but works of AI don't appear to be sufficiently mysterious to raise insurmountable questions about their entitlement to copyright.[3]

Our thesis is that AI systems[4] face--and should face--no more difficulty than traditional computer programs in obtaining copyright protection. Since the world copyright community is in general agreement that computer programs are entitled to copyright, AI systems should enjoy the same status. We do not suggest that all AI systems are copyrightable; they must meet the requirements that any copyrightable work must meet, e.g., be original works of authorship. But the mere fact that a work is an AI system should not disenfranchise it from the world of copyright.

II. What is Artificial Intelligence?

We don't propose to define what artificial intelligence is or to explain it (certainly not for the experts at this conference), but only to state a basic premise for our discussion of AI copyright protection. Our description of various types of artificial intelligence is not authoritative or complete; it provides only a basic outline of the AI characteristics relevant to copyright.[5]

The WIPO Program Notes[6] provide descriptions of three categories of AI: expert systems, perception systems and natural-language systems. We use these categories as AI models for our analysis of copyrightability.[7]

A. Expert Systems

Expert systems, described by WIPO as the most important existing practical application of AI, are computer programs for solving problems in specialized fields of knowledge. Expert systems may be used to diagnose medical conditions and recommend treatment, to diagnose problems in machinery, to engage in financial analysis, to determine geological conditions, and to solve problems in a number of other specialized areas.[8]

There are two main components to any expert system: a knowledge base and an inference engine. The knowledge base contains the expertise of one or more specialists in a particular field ("domain experts"); the expertise typically is expressed as rules in the form of "if-then" statements.[9] For example, one rule in a knowledge base for an expert system that engages in legal reasoning may be: "If the plaintiff was negligent in the use of the product, the theory of contributory negligence applies."[10] Since the domain expert is not likely to have expertise in computer programming, the knowledge base is likely to have expertise in computer programming, the knowledge base is created by a knowledge engineer, someone with computer programming skills who interviews the domain experts and translates their expertise into a language that the computer can understand.[11]

However, an expert system requires more than a collection of specialized knowledge or rules. It requires also the ability to reason--to apply its knowledge to a particular situation. That is the mission of the inference engine, the second component of the AI system. The inference engine contains general problem-solving knowledge that permits it to decide logically how to apply the specialized rules to the facts the user supplies, in order to solve the problem the user poses.[12]

Ready-made expert systems are not unusual. They consist of an inference engine and support tools such as a knowledge base editor[13] and an explanation facility,[14] and theoretically are able to solve problems in any number of areas after the appropriate knowledge base for the area is created and input.[15]

An expert system might be considered a hybrid of a computer program and database.[16] But, as with computer programs generally, the presence of data (or a "database") as a component needn't change for us its essential nature as a computer program--albeit a sophisticated one.

B. Perception Systems

Perception systems (e.g., "computer vision") are systems that permit a computer to "perceive" the world, typically by providing the computer with a "sense" of "sight" or "hearing." An optical character recognition system (OCR), which permits the computer to "read" printed text, is a well-known example of computer vision.[17] In a sense, the more advanced OCR's are expert systems that contain a number of computerized "experts" such as topological experts (e.g., loop, concavity and line-segment detectors, which recognize various aspects of letters) and word-context experts (which have knowledge of the possible words in a given language), as well as an "expert manager" programmed with knowledge as to which expert to use in a particular situation.[18]

There are also computer vision systems that "see" and "recognize" objects in the real world, functions that require an ability to discern shapes and edges of objects that is more advanced than that of the OCR. In addition to "seeing" the object, the computer vision system must also "recognize" it, typically by comparing what it sees to the object models it stores in its knowledge base. The recent conflict in the Persian Gulf provided striking examples of military applications of computer vision systems.[19]

Speech recognition systems do for hearing what computer vision systems do for sight. As with computer vision, speech recognition systems are based on the principle of pattern recognition, requiring a knowledge base not only of words, but also of elements such as phonemes, syntax and semantics.[20]

C. Natural-Language Systems

As we've noted, many perception systems have elements that require, at some level, an understanding of human language. Natural language itself is a promising field in AI, within and outside the context of perception systems. Programs are being developed to translate from one language to another, or to prepare short abstracts of lengthy texts. A program may also have a natural language ability simply to permit a lay user to communicate easily with the computer.

A natural language program must understand the meanings of words, which requires creating a dictionary database (or utilizing a dictionary already available in machine-readable form). However, that is only the beginning. Since words have different meanings in different grammatical and textual contexts, the system must apply the AI technology of semantic analysis and must understand the rules of syntax. A technique called pragmatic analysis utilizes knowledge about the real world to assist the program in making choices about the meaning of words and sentences.[21] A natural language system therefore is very much a knowledge-based system, requiring that the computer have "knowedge" of the rules of language, the meanings of words and the world in general.

III. AI From the Perspective of Copyright

From what we've described, it is apparent that, to determine its copyrightability, a typical AI system can be viewed as consisting of one or more computer programs and one or more databases. The United States copyright statute provides an example of an approach that is helpful in the inquiry.

Under U.S. copyright law, "computer program" has a particular meaning. The Copyright Act defines a computer program as:

> "a set of statements or instructions to be use directly or indirectly in a computer in order to bring about a certain result."[22]

The Draft Model Copyright Law considered by WIPO proposes a similar definition.[23]

Does that definition cover elements of an AI system--e.g., the inference engine and the knowledge base of an expert system? The inference engine appears to be what we traditionally consider to be a computer program: a set of instructions for the computer to execute.[24] The knowledge base of an expert system may be a more complex work, sharing characteristics of both a computer program and a database.

The knowledge base appears to fit within the definition of "computer program" in the Copyright Act since it typically contains statements, and those statements are used directly or indirectly in the computer to bring about a result (e.g., predicting a determination of liability in a product liability case). As we described above, the knowledge base is typically made up of a series of rules in the form of "if-then" statements. Each rule leads either to a conclusion or the application of another rule.[25] In this matter the rules contained in the knowledge base control the flow of the program's execution in the same fashion as the statements of a computer program.[26]

The knowledge base--as well as other AI elements discussed above, such as the stored object models of a perception system or the dictionary of a natural language program--can also be thought of as another type of literary work: a database or other compilation. One definition of a "compilation," similar to that in many copyright laws is:

> "a work formed by the collection and assembling of preexisting materials or of data that are selected, coordinated, or arranged in such a way that the resulting work as a whole constitutes an original work of authorship. The term 'compilation' includes collective works."[27]

A relevant non-legal definition of a "database" would be:

> "A collection of related information about a subject organized in a useful manner that provides a base or foundation for procedures such as retrieving information, drawing conclusions, and making decisions. Any collection of information that serves these purposes qualifies as a database, even if the information is not stored on a computer..."[28]

An expert system's knowledge base is a collection of related information about a particular subject; it is organized in a useful manner; and it provides a foundation for procedures such as drawing conclusions and making decisions. The information included in the knowledge base is selected and coordinated from the entire realm of human knowledge on a given subject.[29] Consequently a knowledge base presents the requisite elements of selection, coordination or arrangement that would make it copyrightable as a database compilation.[30] The natural language system's dictionary and the perception system's library of stored object models similarly qualify as databases.

Although databases are not defined or expressly mentioned in the U.S. Copyright Act, the definition of "compilation" quoted above clearly includes databases ("collection and assembling of...data...," etc.) As CONTU noted, "[t]he unauthorized taking of of substantial segments of a copyrighted data base should be considered infringing, consistent with the case law developed from infringement of copyright in various forms of directories."[31]

The United States Supreme Court's recent decision in Feist Publications Inc. v. Rural Telephone Service Co.[32] places particular emphasis on the statutory words "selected, coordinated, or arranged" as the indicia for finding original expression that is eligible for copyright protection in a compilation. But, the Court made it clear that, although there are some few works (such as the telephone directory white pages at issue in that case) that will not pass muster, the level of creativity required in a particular selection, coordination or arrangement of factual material is quite low.[33]

Under Feist, a perception system's library of stored object models or a natural language system's dictionary also fit comfortably into the statutory definition of a copyrightable compilation and most likely contain sufficient originality to be protectible as a database. Take, for example, a hypothetical system designed to recognize geometric shapes. The compilation is protected as a whole, although a single circle or square stored in the program would not be entitled, in and of itself, to copyright protection. The single circle or square most likely would not be considered to have sufficient orginality to qualify for copyright, but the collection of geometric figures selected by the programmer for inclusion in the program may qualify for protection as a compilation. Similarly, the dictionary in a natural language program may qualify for protection as a compilation, just as a more traditional dictionary would.

Ultimately, however, it probably is not necessary to ask whether these components of AI programs are entitled to copyright protection as database compilations--for at least a few reasons. First, just as many other copyrighted works (e.g., motion pictures) include elements that can be classified in many categories, a computer program frequently includes elements of a database and elements of works that can be classified in other categories. For example, computer programs include audiovisual elements such as screen displays.[34]

Second, in determining copyrightability it may not be necessary (or appropriate) to examine separately--apart from the computer program itself--the individual AI elements such as knowledge base of an expert system, the stored object models of a perception system, or the dictionary in a natural language system.[35] In a similar context (analyzing the copyrightability of an audiovisual work), the U.S. Court of Appeals for the District of Columbia has emphasized the Copyright Act's "apparent recognition that the whole--the 'series of related images'--may be greater than the sum of its several or stationary parts."[36]

Third, as noted above,[37] elements of AI programs such as the knowledge base of an expert system probably qualify on their own merits for protection as computer programs. Whether they also can be considered database compilations would therefore be irrelevant.

AI systems therefore fall within categories of protected works already recognized in the copyright law, even though, as in the United States, the law may not specifically designate "computer programs" or "databases" as separate categories. Again as in the United States, it is common for "computer programs" and "databases" to be protected as "literary works." Literary works are one of the seven categories of "works of authorship" in which copyright protection subsists under the U.S. Copyright Act,[38] and are defined as:

> "works, other than audiovisual works, expressed in words, numbers, or other verbal or numerical symbols or indicia, regardless of the nature of the material objects, such as books, periodicals, manuscripts, phonorecords, films, tapes, disks or cards, in which they are embodied."[39]

By its own terms, this broad definition would include AI systems and traditional computer programs, which are clearly works "expressed in words, numbers, or other verbal or numerical symbols of indicia," and typically embodied in media such as tapes or disks. Indeed, the legislative history of the 1976 Copyright Act reveals that Congress recognized that the "literary works" embraced by copyright would include both computer programs and databases:

> "The term 'literary works' does not connote any criterion of literary merit or qualitative value: it includes catalogs, directories, and similar factual, reference, or instructional works and compilations of data. It also includes computer data bases, and computer programs to the extent that they incorporate authorship in the programmer's expression of original ideas, as distinguished from the ideas themselves."[40]

As such, AI systems, whether viewed as pure computer programs or as a combination of computer programs and databases, thus qualify for copyright protection under the 1976 Copyright Act as literary works.

United States copyright law is not unusual in its approach to protection of computer programs and data bases. Copyright systems generally offer similar protection to computer programs, usually as literary works;[41] and compilations and databases are typically also protected by copyright.[42] Indeed, it is fair to say that there is now an international consensus that computer programs and databases are copyrightable. For example, the proposed Model Copyright Law considered by WIPO's Committee of Experts would include computer programs among the subject matter of copyright, either as a subcategory of "works expressed in writing" (which, in turn, is a subcategory of "literary and artistic works") or as a separate category (but still a category of "literary and artistic works").[43]

The Committee of Experts agreed that computer programs should be included in the non-exhaustive list of literary and artistic works covered by copyright.[44] Moreover, the experts generally agreed that computer programs are entitled to at least the minimum rights prescribed by the Berne Convention for the Protection of Literary and Artistic Works (Paris Act 1971).[45]

The proposed Model Copyright Law would also protect databases. Section 4 states that "collections of works, of expressions of folklore or of mere facts or data, such as encyclopedias, anthologies and data bases which, by reason of selection and arrangement of their contents, are original," shall also be protected as copyrightable works.[46]

Similarly, the European Community has decided as a matter of community law that computer programs are to be protected by copyright. In its Green Paper on copyright, the EC commission recognized that "at national and international level" there is a "general acknowledgement of the advantages for creators, right holders, users and society as a whole of a 'copyright' solution to the problem of ensuring adequate protection of programs against unauthorized reproduction."[47] The Commission acknowledged that in its member States, both case law and legislation had increasingly recognized the application of copyright to computer programs.[48]

It concluded that a directive should be issued, "explicitly protecting computer programs under copyright law in the broad sense";[49] and such a directive is currently under discussion.[50]

The Commission has rejected notions of protecting computer programs under a regime of neighboring rights or a sui generis regime.[51]

Copyright protection for databases is also under consideration by the EC. The Green Paper[52] addressed the question without coming to any specific conclusions;[53] and a Directive is likely to be promulgated this year to harmonize copyright protection for databases within the Community.[54]

IV. A Few Words About "Systems"

We've been discussing copyright protection for AI "systems," using the terminology that WIPO has used[55]--and terminology common in the field. Of course, copyright protection for an AI system doesn't necessarily mean that the "system" (in the sense of an idea, procedure, process, method of operation, etc.) embodied in the particular computer program and/or database is protected; but let us clarify the issue as framed, because at first blush it would appear to present an insurmountable obstacle to copyrightability. "Systems" (in the sense, again, of an idea, etc.) generally are not entitled to copyright protection.[56] So one might question whether an artificial intelligence "system" would be eligible for copyright.

We have here what a distinguished jurist described in an analogous copyright context as "the one-word-meaning-only fallacy."[57] The mere use of the label "system" to identify what might otherwise be a copyrightable work does not disqualify the work from copyright protection. To consider a "system" ineligible for copyright protection regardless of the context in which the term is used would be ironic since context plays such a major role in AI systems used to recognize the meaning of language. But, more to the point: even though "computer system" is frequently used as a synonym for "computer program," the legally infelicitous phrasing is irrelevant to protecting the expression in the program; and there's no reason why the protection for AI "systems" should be determined any differently.

The provision in United States law that disqualifies "systems" from copyright simply elaborates the traditional dichotomy between ideas, which cannot be protected by copyright, and expression, which can be protected[58]--a principle found in most copyright regimes.[59] Thus, if what is meant by "system" is merely what we usually consider to be the process, procedure or method of operation by which tasks are performed, the simple answer is that copyright offers no protection for such a system. However, WIPO doesn't appear to be using the phrase that way; and we certainly are not. We believe that an "AI system" also contains expression of the process or method, and that--as with computer programs generally--such expression is copyrightable.

The U.S. Congress made this clear when it discussed computer programs in the legislative history of the 1976 copyright law revision:

> "Some concern has been expressed lest copyright in computer programs should extend protection to the methodology or processes adopted by the programmer, rather than merely to the 'writing' expressing his ideas. Section 102(b) is intended, among other things, to make clear that the expression adopted by the programmer is the copyrightable element in a computer program, and that the actual processes or methods embodied in the program are not within the scope of the copyright law."[60]

When asking whether an AI "system" is protectible under copyright law, then, one first must ask what is meant by "system." That is what an appellate court did when a litigant contested the copyright of a "replacement parts numbering system" and argued that, in excluding "systems" from copyright, 102(b) of the U.S. Copyright Act automatically disqualified the plaintiff's work from copyright because the work was called a "system":

> All that the idea/expression dichotomy embodied in 102(b) means in the parts numbering system context is that appellant could not copyright the idea of using numbers to designate replacement parts. Section 102(b) does not answer the question of whether appellant's particular expression of that idea is copyrightable.[61]

The court, in other words, properly looked behind the label for the copyrighted work and required that the work itself be examined to determine whether it was subject to copyright.[62] Similarly, an AI "system" should be eligible for copyright protection in the United States and elsewhere for the particular form of expression of that system, but not for the idea--e.g., the idea of using artificial intelligence to accomplish a particular goal.

The principle can be illustrated by MYCIN, one of the pioneering AI expert systems, that physicians use to diagnose causes of infection and recommend drug treatment.[63]

Copyright clearly would not prevent somebody else merely from creating his or her own AI system independently to fit the same description as MYCIN. But creating a system by copying the expression from MYCIN just as clearly would be an act of copyright infringement--whether the expression copied is the literal expression of statements, instructions, data, etc. or their nonliteral expression, their structure, sequence and organization ("SSO").[64]

V. The "Utilitarian" Nature of AI Systems

AI systems are not produced to be read by the fireside on cold wintry nights. Their purpose is to cause a computer to do something. In this respect they do not differ from traditional computer programs. Dissenters have argued that the utilitarian aspect of computer programs is reason for denying or restricting protection for otherwise copyrightable expression; and we can anticipate that AI will provide opportunity for a reprise of the familiar argument. But the argument must be rejected for AI for the same reasons it's been rejected for computer programs generally.[65] As the United States Supreme Court has said, there is:

> "nothing in the copyright statute to support the argument that the intended use or use in industry of an article eligible for copyright bars or invalidates its registration. We do not read such a limitation in the copyright law."[66]

Nor is an AI system an uncopyrightable "useful article." Rather, it is an intangible work of authorship that may be used in a tangible useful article, the computer.[67]

Even if the limitations on protection for useful articles did apply to AI systems, the system itself (e.g., the knowledge base and inference engine in an expert system) contains protectible expression that can be identified separately from, and exist independently of, its utilitarian aspects.[68] It is the expression in the system that is protected by copyright.

As a court recently has observed in rejecting the "useful article"/ "utilitarian" argument,

> "[i]t does not follow that when an intellectual work achieves the feat of being useful as well as expressive and original, the moment of creative triumph is also a moment of devastating financial loss--because the triumph destroys copyrightability of all expressive elements that would have been protected if only they had not contributed so much to the public interest by helping to make some article useful."[69]

In short, the most that can be said for the "useful article" argument is that the computer is a useful article that is not subject to copyright protection, just as a motion picture displayed by the projector, so too does it protect the expression in the computer program input into the computer. As CONTU noted,

> "[p]rograms should no more be considered machine parts than videotapes should be considered parts of projectors or phonorecords parts of sound reproduction equipment. All three types of works are capable of communicating with humans ... All that copyright protection for programs, videotapes, and phonorecords means is that users may not take the works of others to operate their machines. In each instance, one is always free to make the machine do the same thing as it would if it had the copyrighted work placed in it, but only by one's own creative effort rather than by piracy."[70]

Throughout the world, utilitarian works have long been eligible for copyright.[71]

But the dissenters have failed to cope with that long history of protection, which can be illustrated in the United States by the span of two centuries from its first copyright statute,[72] protecting maps and charts, to the current U.S. copyright law, protecting many utilitarian literary works, such as encyclopedias, code books, directories, fact compilations, dictionaries, and "how-to" manuals.[73]

Indeed, most of the AI systems we have considered would be even less subject to the "utilitarian" argument than are many traditional computer programs. In most cases, a major component of the AI system is a knowledge base that contains information or knowledge. Yes, knowledge is utilitarian, but any argument that its usefulness to the user (or to the reader of an encyclopedia or "how-to" book) bars protection for its expression should not carry far--though it is no more far-fetched than the "useful article" argument for computer programs generally. The fact that the knowledge is placed in a computer's memory and used to solve a problem or perform a task otherwise undertaken by a human being doesn't transform the particular expression of that knowledge into a useful article. The knowledge base is clearly also a copyrightable work.

VI. Copyright vs. "Sui Generis" Protection

Developments in AI may prompt past opponents of copyright for computer programs to revive earlier calls for a sui generis approach. They may contend that AI is more different from conventional literary works than are computer programs generally, and that the argument is stronger for a separate regime custom-tailored to suit future AI systems, i.e., insofar as we can guess at their dimensions.

But the argument carries even less weight for AI than for computer programs in general. There is a very strong consensus in the intellectual property community that copyright is the appropriate vehicle for protection of computer programs and databases,[74] a consensus reflected also in the United States Copyright Act and generally in the copyright systems of other countries.[75] Moreover, international copyright protection for computer programs is secured under the Berne Convention.[76] Any partial dismantling of the international copyright system to adopt a new, non-copyright regime for AI would likely deny AI international protection unless and until a new international convention specifically designed for AI could be adopted. Even if there were such a convention, it is doubtful that it would be subscribed to so universally as Berne, and the result would be that many AI systems would fall into a gap where no protection is available.

Traditional copyright law offers a long history of legislative and judicial interpretation, and generally accepted principles of copyright jurisprudence sufficiently flexible to adapt to new technological developments. To encourage the creation and dissemination of AI for the benefit of society, AI systems should be protected by copyright--rather than being left to the tender mercies of some new, untested substitute--no less than conventional computer programs and other copyrighted works. Copyright is especially appropriate for AI, where rapid technological advances would threaten to make any new statutory scheme obsolete almost as soon as it was enacted. It would be far better to continue the protection of AI under the system of protection that copyright offers, a system that has long proven its ability to adapt.

The sui generis approach, moreover, would offer a convenient excuse--for some who might wish to do so--to water down the scope of protection and make AI systems a poor relation to other works of intellectual property. That's hardly an excuse for granting lesser protection for AI. In fact, one could argue that AI systems have a stronger claim to copyright protection than conventional computer programs, in as much as AI systems typically have a knowledge or data base component more readily accepted as "intellectual property" even by those who have advocated sui generis protection for computer programs based on their perception of them as utilitarian works or "useful articles."[77]

VII. Conclusion: Why AI Systems Should be Protected

Merely to ask whether copyright protection should be recognized for artificial intelligence systems may imply doubt (we have little) as to whether it's recognized now. Computer programs are protected by copyright, and AI systems must be too. There is protectible expression--under present laws--in the components of an AI system (the inference engine and other computer programming, the knowledge base, etc.) and in its whole. Indeed, as we have discussed, in some respects the basis for copyright protection for AI systems may be even more compelling than that for conventional computer programs.

Developing computer programs generally is not inexpensive; and, for some AI systems (by their very nature, knowledge-intensive) it can be even more costly and time-consuming.[78] We doubt that major projects in the AI world, or even their less ambitious cousins, would ever come to fruition if their creators and developers--and investors--believed that others could freely appropriate the fruit of their labors and investment.

As the Copyright Clause of the United States Constitution recognizes, copyright protection is conferred as an incentive "to promote the progress of science and useful arts."[79] It is for that reason that the U.S. Congress is thereby authorized to "secur[e] for limited times to Authors...the exclusive Right to their respective Writings..."[80] As the basis for copyright protection, other legal systems may place greater emphasis on the entitlement of the creator to reward for his or her creation of a work of authorship. But whether the basis is entitlement or incentive, the result is largely the same: without protection for the work many authors would occupy themselves with other tasks that would offer the hope of remuneration. It's true of authors of novels, art and motion pictures--and it's true of authors of artificial intelligence systems.

AI systems promise great value for mankind. But without the exclusive rights of copyright, the production of AI systems will be impaired. If it is, the loss will be not only for those who otherwise would have created AI systems, but for all of us who otherwise would have benefitted.

NOTES

1 World Intellectual Property Organization, WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence; General Information and Provisional Program 3, 7 (Doc. No. SAI/INF/1, November 12, 1990) ("WIPO Program Notes").

2 See, however, R. PENROSE, THE EMPEROR'S NEW MIND (1989), and H. & S. DREYFUS, MIND OVER MACHINE (1986), for analyses that attempt to debunk the notion that AI is or can be truly "intelligent."

3 Other aspects of intellectual property law also offer protection to many AI systems. Although copyright is the most widely-used form of intellectual property protection for software, computer programs sometimes can be protected as trade secrets, See G. Peterson, Trade Secret Rights in Computer Technology, in THE LAW AND BUSINESS OF COMPUTER SOFTWARE 3.03(b) (1990); University Computing Co. v. Lykes-Youngstown Corp., 504 F.2d 518 (5th Cir. 1974); Warrington Assoc., Inc. v. Real Time Engineering Systems, Inc., 522 F. Supp. 367 (N.D. III. 1981); and software related inventions (if not the software itself) may be protected, in appropriate circumstances, by patents. See II C. SHERMAN, H. SANDISON & M. GUREN, COMPUTER SOFTWARE PROTECTION LAW 403.1(c) (1990); In re Iwahashi, 888 F.2d 1370 (Fed. Cir. 1989). In general, however, it is more difficult to meet the requirements for trade secret and patent protection than the requirements for copyright. Because some AI systems may find that copyright is the only form of intellectual property protection available to them, we confine ourselves to examining the copyrightability of AI systems.

4 "System" can be a dangerous word in copyright. See discussion infra, in text at nn. 51-60. However, with a caveat, we defer to WIPO's choice of terminology, see n. 1, supra, and accompanying text. The caveat is that when we refer to an AI "system," we refer to the computer program and/or database in the particular AI application, and not to the method or process of the AI application, or to its procedure, concept, principle, etc.

5 We base the descriptions that follow not upon technical expertise of our own, but upon our review of the literature and discussions with AI specialists.

6 Supra, n. 1, at 3.

7 The WIPO Program Notes do not mention neural networks, still another branch of AI, which take a different approach from expert systems, perception systems and natural-language systems. This may be because, as another panelist at this symposium has noted, neural network systems are still in their infancy and are "perhaps to trivial to warrant consideration of the concept of artificial intellect." A. Johnson-Laird, Neural Networks: The Next Intellectual Property Nightmare?, THE COMPUTER LAWYER 7,14 (March 1990). Johnson-Laird anticipates that within the next 10 years, neural networks may be developed that approximate the intelligence of a bumblebee. Id. at 14. Although, as Johnson-Laird observes, this is "no mean feat,"id., it suggests possibly the limited practical benefits that can be expected from neural networks in the foreseeable future. In any event, we follow the WIPO Program Notes in not including them in our consideration at this time.

8 F. Hayes-Roth, Expert Systems, in ENCYCLOPEDIA OR ARTIFICIAL INTELLIGENCE 287, 288 (1990). See generally D. WATERMAN, A GUIDE TO EXPERT SYSTEMS 244-299 (1986) ("WATERMAN"), for a catalog of expert systems, with brief descriptions of each system.

9 WATERMAN at 20-21 and 63-69. Alternatively, the knowledge base may be organized in "semantic nets" which are based on relationships between objects, concepts, or events (e.g. the relationship expressed in the statement, "The Queen Mary is an ocean liner"), or in "frames" which associate features or attributes with nodes representing concepts or objects. Id. at 70-77.

10 Id. at 16. Parenthetically, we note that such a legal rule would be applicable presumably in a jurisdiction that recognizes contributory negligence, but would have to be modified in those jurisdictions following the trend to adopt comparative negligence. See, e.g., Liv. Yellow Cab Company of California, 13 Cal.3d 804, 532 P.2d 1226 (1975).

11 WATERMAN at 9.

12 WATERMAN at 18-19; M. COVINGTON AND D. DOWNING, DICTIONARY OF COMPUTER TERMS 120 (1989)

13 A knowledge base editor assists in loading information into the knowledge base. WATERMAN at 92-93.

14 An explanation facility explains to the user how the expert system has reached its particular conclusions. WATERMAN at 39-91. See W.R. Swartout, Explanation, in ENCYCLOPEDIA OF ARTIFICIAL INTELLIGENCE, supra, n.8, at 298-300.

15 F. Hayes-Roth, supra, n. 8, at 294-295; WATERMAN at 83.

16 It is probably only a slight oversimplication to consider the inference engine as a rather traditional form of computer program and the knowledge base as sharing characteristics of both a computer program and database. We will revisit this characterization later to consider its copyright implications. See discussion infra in text at nn. 25-50.

17 R. KURZWEIL, THE AGE OF INTELLIGENT MACHINES ("KURZWEIL") 272-275 (1990)

18 KURZWEIL at 238-247, 272-274.

19 For a general discussion of computer vision, see id. at 247-262.

20 Id. at 263-270

21 Id. at 303-312

22 Section 101 of the Copyright Act of 1976, 17 U.S.C. 101. The definition of "computer program" was added to 101 in 1980 as part of an amendment recommended by the National Commission on New Technological Uses of Copyrighted Works (CONTU). The definition is taken from the FINAL REPORT OF THE NATIONAL COMMISSION ON NEW TECHNOLOGICAL USES OF COPYRIGHTED WORKS (1979) ["CONTU FINAL REPORT"] at 12.

It is important to note, however, that the definition of a copyrighted work is not the definition of the scope of its protection. To illustrate, a motion picture is not the definition of the scope of its protection. To illustrate, a motion picture is defined in the U.S. statute as, essentially, a work that consists of "a series of related images" which impart a sense of motion, "together with accompanying sounds, if any" (17 U.S.C. 101); but the copyright on a motion picture does not protect merely the "images" and the "sounds." It also protects, for example, the detailed plot of the motion picture. Similarly, see discussion at n. 62, infra, as to the scope of copyright protection as a computer program, embracing not merely literal text of the work's "statements" or "instructions" but also its nonliteral structure, sequence and organization ("SSO").

23 The WIPO International Bureau has proposed the following definition:

> "A 'computer program' is a set of instructions expressed in words, codes, schemes or in any other form, which is capable, when incorporated in a machine-readable medium, of causing a 'computer'--an electronic or similar device having information-processing capabilities--to perform or achieve a particular task or result."

World Intellectual Property Organization, Committee of Experts on Model Provisions for Legislation in the Field of Copyright, First Session, Geneva, February 20 to March 3, 1989, DRAFT MODEL PROVISIONS FOR LEGISLATION IN THE FIELD OF COPYRIGHT, Memorandum Prepared by the International Bureau of the World Intellectual Property Organization (WIPO), Doc. No. CE/MPC/I/2-II (October 20, 1988), sec. 1(vii).

24 See M. COVINGTON AND D. DOWNING, DICTIONARY OF COMPUTER TERMS 120 (1989).

25 See text at n. 9, supra.

26 See text at nn. 9-10, supra.

27 Section 101 of the U.S. Copyright Act of 1976, 17 U.S.C. 101.

28 Another definition has been proposed in the Commission of the European Communities in its GREEN PAPER ON COPYRIGHT AND THE CHALLENGE OF TECHNOLOGY--COPYRIGHT ISSUES REQUIRING IMMEDIATE ACTION (Commission Document COM (88) 172 final, Brussels, of June 7, 1988), 6.1.1:

> "The term 'data base' is used in this chapter to mean a collection of information stored and accessed by electronic means. It may be a collection of full-text material, that is to say, existing copyright works, in which case an analogy might be made between the data base and a generalized or specialized library. It may be a compilation of extracts of works, similar to an anthology or a documentation centre, from which relevant parts of works may be obtained. It may be a collection of material which is in the public domain, such as lists of names and addresses, prices, reference numbers. There is here a similarity with catalogues, timetables, price lists and other such reference material in printed form. Lastly, it may consist of the electronic publishing of a single but voluminous work, such as an encyclopedia."

Subsequently, in the FOLLOW-UP TO THE GREEN PAPER (COM/584/90-FINAL, Dec. 5, 1990) 6.2.2.(2), the Commission reported on responses to its questionnaire on the protection of databases:

> "As regards a definition of 'database', several participants proposed a broad definition which includes the following elements:
>
> a) collection, organization and storage of data;
>
> b) information in a digital form in which it can be processed by means of a computer."

It is self evident that the knowledge bases of various AI systems would fit within this definition.

29 Moreover, that information is expressed in particular statements that would be subject to copyright protection.

30 See the discussion which follows immediately hereafter.

31 CONTU FINAL REPORT AT 42.

32 499 U.S.___, 113 L.Ed.2d 358, 59 U.S.L.W. 4251, 18 U.S.P.Q.2d 1275,18 Med.L.Rptr. 1889, No. 89-1909 (March 27, 1991)

33 Id., 18 U.S.P.Q.2d at 1283-1285, slip op. at 17, 21-23.

34 "Copyright protection applies to the user interface, or overall structure and organization of a computer program, including its audiovisual displays, or screen 'look and feel.'"

Telemarketing Resources v. Symantec Corp., 12 U.S.P.Q.2d 1991, 1993 (N.D. Cal. 1989). See also Lotus Development Corp. v. Paperback Software Int'l, 740 F. Supp. 37, 79-80 (D. Mass 1990). The U.S. Copyright Office expresses the same view, in its Notice of Decision concerning Registration and Deposit of Computer Screen Displays, 53 Fed. Reg. 21817 (June 10, 1988):

> "...[A] copyrightable expression owned by the same claimant and embodied in a computer program, or first published as a unit with a computer program, including computer screen displays, is considered a single work... Where a work contains different kinds of authorship, the registration class will be determined on the basis of which authorship predominates."

35 See Allen-Myland, Inc. v. International Business Machines Corp., 746 F. supp. 520, 531-532 (E.D.Pa. 1990), where the court rejected the counterclaim defendant's contention that a portion of the copyright owner's microcode consisting largely of a list of parts was not copyrightable because it lacked sufficient originality. The court pointed out that the copyright owner's 3090 system "cannot function properly" without that portion of the microcode, and rejected the attempt to break up the program into its component parts, concluding that "the work must be reviewed as a whole, not just reviewed or analyzed part by part." (Quoting M. Kramer Mfg. Co., Inc. v. Andrews, 783 F2d 431, 439 (4th Cir. 1986).) Database compilations are commonly elements of a computer program. But even if not, a computer program should not be denied

protection for programming elements (especially those otherwise protectible, such as data compilations) merely because the elements might not appear to fit comfortably into a traditional definition of computer programs. Indeed, the statutory definition of "computer programs" is not the provision of the Copyright Act that gives them protection; rather, computer programs are copyrightable as literary works. See Section 106 of the Copyright Act of 1976, 17 U.S.C. 106, the definition of "literary works" in section 101, 17 U.S.C. Sec. 101, and the discussion of "literary works" in text at nn. 38-40, infra.

36 Atari Games Corp. v. Oman, 888 F.2d 878, 881-882 (D.C. Cir. 1989). See also n. 22, supra.

37 See text at nn. 25-26, supra.

38 See Section 102(a) of the Copyright Act of 1976, 17 U.S.C. Sec. 102(a).

39 Section 101 of the Copyright Act of 1976, 17 U.S.C. Sec. 101.

40 H.R. REP. NO. 1476, 94th Cong., 2d Sess. 54 (1976) (emphasis supplied)

41 See, e.g., Copyright Act, Sec. 10(1) (definition of literary work) [Australia]; Copyright Statute, Art. 2 [Dominican Republic]; Copyright Statute [Law No. 57-298 on Literary and Artistic Property (as amended up to July 3, 1985, including Law No. 85-660 of July 3, 1985, Journal Officiel 7495 (1985))], Art. 3 [France]; Copyright Statute, Art. 2(1), item 1 [Germany]; Copyright Act of 1957 (as amended by Copyright Amendment Bill No. XIX of 1984), 2(o) [India]; Copyright Act 1982 [as amended up to September 19, 1987] Arts. 1(g) & 11(k) [Indonesia]; Copyright Law, Sec. 2(1)(xbis) and 10(1)(ix) [Japan]; the Copyright Act 1987, Sec. 7 (definition of literary work) [Singapore]; and Copyright, Designs and Patents Act 1988, Sec. 3(1)(b) [United Kingdom]. Copyright protection for computer software has also been recognized by judicial action in many countries where there is no legislation expressly recognizing it.

For surveys of computer program copyright laws throughout the world, see M. Keplinger, International Protection for Computer Programs, 1991 PACIFIC RIM COMPUTER LAW CONFERENCE IV-3 (Computer Law Association 1991); J. Keustermans & I. Arckens, INTERNATIONAL COMPUTER LAW ch.7 (1988). See also M. Kindermann, The International Copyright of Computer Software, Copyright 201 (April 1988).

42 See, e.g., Copyright Act, Sec. 10(1) (definition of literary work) [Australia]; Law No. 158 of 1961 on Copyright in Literary and Artistic Works (as amended) Sec. 49 [Denmark]; Copyright Statute, Art. 2 [Dominican Republic]; Copyright Act of 1957 (as amended by Copyright Amendment Bill No. XIX of 1984), Sec. 2(o) [India]; Copyright Law Sec. 2(I)(xter), 12 and 12bis [Japan]; The Copyright Act 1987, Sec. 7 (definition of literary work) [Singapore].

43 World Intellectual Property Organization, Committee of Experts on Model Provisions for Legislation in the Field of Copyright, First Session, Geneva, February 20 to March 3, 1989, DRAFT MODEL PROVISIONS FOR LEGISLATION IN THE FIELD OF COPYRIGHT, Memorandum Prepared by the International Bureau of the World Intellectual Property Organization (WIPO), doc. No. CE/MPC/I/2-II (October 20, 1988) Secs. 3(1)(i) and 3(1)(xii).

44 World Intellectual Property Organization, Committee of Experts on Model Provisions for Legislation in the Field of Copyright, Third Session, Geneva, July 2 to 13, 1990, PREPARATORY DOCUMENT, DRAFT MODEL LAW ON COPYRIGHT, Memorandum Prepared by the International Bureau of the World Intellectual Property Organization (WIPO), doc. No. CE/MPC/III/2 (March 30, 1990), 144. See also World Intellectual Property Organization, Committee of Experts on Model Provisions for Legislation in the Field of Copyright, Third Session, Geneva, July 2 to 13, 1990, REPORT ADOPTED BY THE COMMITTEE, Doc. No. CE/MPC/III/3 (July 13, 1990), 82.

45 Id. 85.

46 World Intellectual Property Organization, Committee of Experts on Model Provisions for Legislation in the Field of Copyright, First Session, Geneva, February 20 to March 3, 1989, DRAFT MODEL PROVISIONS FOR LEGISLATION IN THE FIELD OF COPYRIGHT, Memorandum Prepared by the International Bureau of the World Intellectual Property Organization (WIPO), Doc. No. CE/MPC/I/2-II (October 20, 1988) Sec. 4(1)(ii) (emphasis supplied).

47 Commission of the European Communities, GREEN PAPER ON COPYRIGHT AND THE CHALLENGE OF TECHNOLOGY--COPYRIGHT ISSUES REQUIRING IMMEDIATE ACTION (Commission Document COM (88) 172 final, Brussels, of June 7, 1988), 5.3.4.

48 Id. 5.3.7 % 5.3.8. See n. 41, supra.

49 Id, 5.6.2. See id., 5.8.1.

50 See, e.g., Commission of the European Communities, PROPOSAL FOR A COUNCIL DIRECTIVE ON THE LEGAL PROTECTION OF COMPUTER PROGRAMS (Commission Document COM (88) 816 final--SYN 183, of January 5, 1989); Commission of the European Communities, FOLLOW-UP TO THE GREEN PAPER, (Commission document COM/58490-FINAL, Brussels, December 5, 1990); Council of the European Communities, COMMON POSITION ADOPTED BY THE COUNCIL ON 13 DECEMBER 1990 WITH A VIEW TO THE ADOPTION OF A DIRECTIVE ON THE LEGAL PROTECTION OF COMPUTER PROGRAMS (Document No. 10652/1/90, Brussels, 14 December 1990) ("COMMON POSITION").

51 FOLLOW-UP TO THE GREEN PAPER, supra n. 50, 5.2.2.(b).

52 Supra n. 47.

53 See FOLLOW-UP TO THE GREEN PAPER, supra n. 50, 6.1.2. ("The conclusions of this chapter of the Green Paper were left relatively open ended, with no firm indication being given of specific action by the Commission in view of the rapid development of this new sector.")

54 See id., 6.2.2. and 6.3.

55 See text at n. 1, supra.

56 For example, the United States copyright statute provides that copyright protection does not extend "to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied..." Section 102(b) of the Copyright Act of 1976, 17 U.S.C. Sec. 102(b).

57 In an era when the U.S. copyright jurisprudence was riddled and saddled with formalities relating to "investive" publication (obtaining copyright by first public distribution of copies with proper notice affixed) and "divestive" (forfeiting copyright by distribution without proper notice), Judge Jerome Frank observed:

> "In deciding whether certain acts constitute 'publication',... numerous conflicting cases ... by their holdings, though not in their stated rationale, raise more than a suspicion that the term "publication" is clouded by semantic confusion where the term is defined for different purposes, and that we have here an illustration of the one-word-meaning-only fallacy... It is, however, perfectly clear that the word 'publication' does not have the same legal meaning in all contexts... [T]he courts apply different tests of publication... In each case the courts appear so to treat the concept of 'publication' as to prevent piracy."

American Visuals Corp. v. Holland, 239 F.2d 740, 742-744 (2d Cir. 1956)

58 See n. 56, supra.

59 See World Intellectual Property Organization, BACKGROUND READING MATERIAL ON INTELLECTUAL PROPERTY 209 (1988). Thus, for example, under French law, ideas, concepts and methods are not protected by copyright. R. Plaisant, France, 2[1][b][i], in P. Geller & M. Nimmer, INTERNATIONAL COPYRIGHT (1990).

60 H.R. REP. NO. 1476, 94th Cong., 2d Sess. 57 (1976)

61 Toro Company v. R&R Products Co., 787 F.2d 1208 (8th Cir. 1986). See also 1 M. & D. NIMMER, NIMMER ON COPYRIGHT 2.03[D], which takes the position that it would be a misreading of 102(b)

> "to interpret it to deny copyright protection to 'the expression' of a work, even if that work happens to consist of an 'idea, procedure, process, etc.' Thus, if a given 'procedure' is reduced to written form, this will constitute a protectible work of authorship so as to preclude the unlicensed copying of 'the expression' of the procedure, even if the procedure per se constitutes an unprotectible 'idea." Therefore, although 102(b) denies that copyright may 'extend to' an 'idea, procedure, process, etc.,' as contained in a given work it does not deny copyright to a work merely because that work consists of an 'idea, procedure, process, etc.'"

62 See also Apple Computer, inc. v. Franklin Computer Corp., 714 F.2d 1240 (3d Cir. 1983), cert. dismissed, 464 U.S. 1033 (1984) where the court upheld the copyrightability of operating system programs despite the defendant's objection that an operating system is either "process, "system," or method of operation."

63 One commentator's description of MYCIN is as follows:

> "MYCIN assists physicians in the selection of appropriate antimicrobial therapy for hospital patients with bacteremia, meningitis, and cystitis infections. The system diagnoses the cause of the infection (e.g., the identity of the infecting organism is pseudomonas) using knowledge, relating infecting organisms with patient history, symptoms, and laboratory test results. The system recommends drug treatment (type and dosage) according to procedures followed by physicians experienced in infectious disease therapy. MYCIN is a rule-based system employing a backward chaining control scheme. It includes mechanisms for performing certainty calculations and providing explanations of the system's reasoning process. MYCIN is implemented in LISP."

WATERMAN, supra n. 8, at 283.

64 Whelan Associates, Inc. v. Jaslow Dental Laboratory, Inc., 797 F.2d 1222 (3d Cir. 1986), cert. denied, 479 U.S. 1031 (1987); Johnson Controls, Inc. v. Phoenix Control Systems, Inc., 886 F.2d 1173 (9th Cir. 1989); SAS Institute, Inc. v. S&H Computer Systems, Inc. 605 F. Supp. 816 (M.D. Tenn. 1985); E.F. Johnson Co. v. Uniden Corp. of America, 623 F. Supp. 1485 (D. Minn. 1985); Broderbund Software Inc. v. Unison World, Inc., 648 F. Supp. 1127 (N.D. Cal. 1986); Dynamic Solutions, Inc. v. Planning & Control, Inc. [1987] Copyright L. Dec. (CCH) 26,062 (S.D.N.Y. 1987); Pearl Systems, Inc. v. Competition Electronics, Inc., 8 U/S/P/Q/2d 1520 (S.D. Fla. 1988); Soft Computer Consultants, Inc. v. Lalehzarzadeh, [1989] Copyright L. Dec. (CCH) 26,403 (E/D/N/Y/ 1988); Manufacturers Technologies, Inc. v. CAMS, Inc., 706 F. Supp. 984 (D. Conn. 1989); Lotus Development Corp. v. Paperback Software Int'l, 740 F. Supp. 37 (d. Mass. 1990); Atari Game Corp v. Nintendo of America, Inc., Nos. C-88-4805-FMS, C-89-0027-FMS, C-89-0824-FMS (N.D. Cal. Mar. 27, 1991); Customs Service Decision 90-40 (Jan. 10, 1990) File: HQ 732291 CPR-3 CO:R:C:V 732291 SO. 24 Cust. B. & Dec. No. 14, p. 28 [1990] Guide to Computer Law (CCH) 60,212 (Apr. 4, 1990). See also, Q-Co. Industries, Inc. v. Hoffman, 625 F. Supp. 608 (S.D.N.Y. 1985); Digital Communications Associates, Inc. v. Softklone Distributing Corp., 659 F. Supp 449 (N.D.Ga. 1987); Healthcare Affiliated Services, Inc. v. Lippany, 701 F. Supp. 1142 (W.D. Pa. 1988); Telemarketing Resources v. Symantec Corp., 12 U.S.P.Q. 2d 1991 (N.D. Cal. 1989); Bull HN Information Systems, Inc. v. American Express Bank Limited, [1990] Copyright L. Dec. (CCH) 26,555 (S.D.N.Y. 1990).

But compare Plains Cotton Cooperative Assn. v. Goodpasture Computer Service, Inc., 807 F.2d 1256 (5th Cir.), cert. denied, 484 U.S. 821 (1987); Synercom Technology, Inc. v. University Computing Co., 462 F. Supp. 1003 (N.D. Tex. 1978).

65 See, e.g., Apple Computer, Inc. v. Franklin Computer Corp., 714 F. 2d 1240, 1249 (3d Cir. 1983), cert. dismissed, 464 U.S. 1033 (1984); Williams Electronics, Inc. v. Artic International, Inc.., 685 F.2d 870, 874, 876 (3d Cir. 1982); Lotus Development Corp. v. Paperback Software Int'l, 740 F. Supp. 37, 71, 72 (D. Mass. 1990); Apple Computer , Inc. v. Formula International, Inc., 562 F. Supp. 775 (C.D. Cal. 1983), aff'd 725 F.2d 521 (9th Cir. 1984); Midway Mfg. Co. v. Artic International, Inc. 547 F. Supp. 999 (N.D. Ill. 1982), aff'd, 704 F.2d 1009 (7th Cir. 1983); E.F. Johnson Co. v. Uniden Corp. of America, 623 F. Supp. 1498 (D. Minn. 1985); NEC Corp v. Intel Corp., 645 F. Supp. 590, 595 (N.D. Cal. 1986), vacated on grounds of judge's recusal, see 835 F.2d 1546 (9th Cir. 1988). See also discussion of "fallacies and fables about 'useful articles'" in M. Goldberg and J. Burleigh, Copyright Protection for Computer Programs: Is the Sky Falling?, AIPLA Q.J. 294, 319-322 (1989).

66 Mazer v. Stein, 347 U.S. 201, 218 (1954). The Court referred to "registration" because registration was required under the formalities of earlier copyright law.

67 In any event, under United States copyright law the restrictions on protection for "useful articles" are confined to the special category of pictorial, graphic and sculptural works--a category into which computer programs hardly fit. See 17 U.S.C. Sec. 102(a)(5) and the definition of "pictorial, graphic and sculptural works" in 17 U.S.C. Sec. 101 ("...Such works shall include works of artistic craftsmanship insofar as their form but not their mechanical or utilitarian aspects are concerned; the design of a useful article, as defined in this section, shall be considered a pictorial, graphic, or sculptural work only if, and only to the extent that, such design incorporates pictorial, graphic, or sculptural features that can be identified separately from, and are capable of existing independently of, the utilitarian aspects of the article.") See also E.F. Johnson Co. v. Uniden Corp. of America, supra, n. 63, at 1498; I.P. GOLDSTEIN, COPYRIGHT, 2.5.3, n.66 (1989)

68 See n. 67, supra.

69 Lotus Development Corp. v. Paperback Software Int'l, 740 F. Supp. 37, 57 (D. Mass. 1990). Judge Keeton correctly concluded that "the mere fact that an intellectual work is useful or functional--be it a dictionary, directory, map, book of meaningless code words, or computer program--does not mean that none of the elements of the work can be copyrightable." Id. at 58.

70 CONTU FINAL REPORT at 21.

71 Such protection is not confined to the United States. The WIPO GUIDE TO THE BERNE CONVENTION (1978) notes that a work "may be produced for purely educational purposes or with a merely utilitarian or commercial aim, without this making any difference to the protection it enjoys." Id., 2.1

72 Act of May 31, 1790, 1 Stat. 124.

73 See text at n. 40, supra.

74 See Samuelson, Survey on the Patent/Copyright Interface for Computer Programs, 17 AIPLA Q.J. 256 (1989), in which Samuelson, an advocate of a sui generis approach, reports the results of a survey she undertook of intellectual property lawyers. They largely favored continuing to work within the copyright and patent systems to achieve the proper balance of intellectual property protection for computer programs, rather than to adopt some sort of sui generis system of protection for programs. Id. at 281. She reports this preference as '[t]he strongest consensus of all among the survey respondents." Id. at 260.

75 See discussion, supra, at nn. 41-42.

76 See discussion, supra, AT NN. 43-45.

77 In our discussion in the text at nn. 65-73, supra, we have responded to the "useful article"/"utilitarian" argument.

78 For example, one ambitious AI project, CYC, attempts to teach a computer "common sense" by programming tens of millions of items of knowledge (in 5 billion bytes) over a ten-year period. CYC has a gigantic knowledge base and 25 inference engines. If successful, CYC promises to do what AI finds it most difficult to do: learn the common human knowledge that would enable it to know what every human knows, rather than simply learn the expertise of a human expert in a particular field. See G. Rifkin, Packing Some Sense into Computers, COMPUTERWORLD 22 (October 15, 1990); Artificial Intelligence; Child's Play, THE ECONOMIST 80 (January 12, 1991).

79 U.S. CONST., Art. I, Sec. 8, Cl. 8.

80 Id.

PROTECTING EXPERT SYSTEMS, IN PARTICULAR EXPERT SYSTEM KNOWLEDGE: A CHALLENGE FOR LAWYERS

by

Jaap H. Spoor
Vrije Universiteit and Trenité van Doorne
Amsterdam, Netherlands

I. Introduction

Expert systems confront the law with new problems as well as opportunities. These range from their application in the legal domain in the form of legal expert systems to questions concerning liability for system failures and protection of systems or knowledge against copying or imitation. Hereunder, I will limit myself to some aspects of the protection issue.

Waterman describes expert systems as "sophisticated computer programs that manipulate knowledge to solve problems efficiently and effectively in a narrow problem area."[1] This definition is as broad as it is simple, and thus leaves room for the wide variety of expert systems that already have been developed or are being built now. For the domain certainly is wide, and so is the field of study as to how expert systems can be protected by intellectual property law.

Standard type expert systems will at least consist of a knowledge base and an inference engine. The knowledge, supplied by (or often rather extracted from) an expert in a given domain is translated by a knowledge engineer into formalized rules, frames or other structures and as such implemented in the system's knowledge base. The inference engine--consisting of software--serves as a tool for applying the knowledge to problems in the given domain, hopefully in order to solve them. Further software may come in useful, such as a knowledge editor which assists in implementing the knowledge, an explanation facility which helps understanding the answers, and a report generator.

As already pointed out this description merely refers to what can be seen as a standard expert system. However, in practice expert systems vary a good deal in many respects. In size, for instance. HEMA, a chain of department stores in the Netherlands, is testing its Paint Adviser, a modest system which informs customers what paint, brushes and other tools they need for and how to proceed with the painting job they have in mind.[2] It runs on an IBM AT or compatible. At the other end of the scale, Micro-electronics and Computer Technology Corp. (MCC) in Austin, Texas, is steadily building a Large Common Sense Knowledge Base called Cyc. A team of twelve has been on the job since 1984 and will continue at least until 1995, trying to implement in Cyc's

---

1 Donald A. Waterman, A Guide to Expert Systems. Reading, Mass. (Addison-Wesley) 1986, p. xvii.

2 R. van der Spek, [3] Kennissystemen, November 1989, p. 5.

database, by means of some 100 million inferences, the everyday knowledge of the world we all have or make believe we have. Cyc is meant to be consulted by other computer programs whenever they would need such everyday knowledge. It clearly does not run on a PC.[3]

Expert systems not only vary in size but in many other respects as well. The kind of knowledge they work with depends on the field of application and so does the method of representing it. Some systems use standard or custom built software shells, in others the inferencing mechanism and the knowledge are firmly interwoven. Different tasks require different software languages which in return lead to differences of approach. Originality and thereby copyrightability may be influenced by the chosen form of knowledge representation.

We will come back to some of these differences later. For the moment we will simply consider expert systems as consisting of two elements: software and knowledge. I will first briefly deal with the software and then in more detail with the knowledge.

II. Protecting Expert System Software

Software is protected by several regimes, including patents and copyright. In some countries, inventive software may as such be patented. However, most countries limit protection to inventions of a technical nature. While technical software can therefore be patented in many countries, administrative software usually cannot, as is for instance demonstrated by current EPO practice.

In its well-known Vicom decision, the EPO granted a patent for image processing software. The images were considered as physical entities, although they were stored as electric signals.[4] However, inventions with respect to document abstracting and retrieving,[5] word processing,[6] and spelling checking[7] were merely seen as mental acts and thereby considered unpatentable. Besides, patents are for inventions only, and novel inventions at that. Although expert system development is always costly, it is not necessarily inventive. Moreover, granting always take years. Patents therefore only offer protection in a limited number of instances. For more generally available protection we must turn towards copyright.

---

3 Doug Millison, in Computable (Dutch edition) 1 March 1991, p. 35.

4 EPO Technical Board of Appeal 15 July 1986, EIPR 1987 p. D 100.

5 EPO Technical Board of Appeal 5 October 1988, Official Journal EPO 1-2/1990, p. 12.

6 EPO Technical Board of Appeal 14 February 1989, Official Journal EPO 9/1990, p. 384.

7 EPO Technical Board of Appeal 14 March 1989.

At present copyright is certainly the main instrument for software protection; expert system software is no exception. Developing expert system software requires considerable creativity, and newly developed inference engines or knowledge editors will therefore be considered original and protected by copyright under all but the most exacting of national copyright laws. Views are more likely to differ if only certain aspects of the system are copied, such as the problem analysis or the user interface. Still, these aspects need not be discussed here, since they differ in no way for expert system software as compared to other software products.

III. Protecting Expert System Knowledge

Expert systems process knowledge. Assembling, assessing and implementing the relevant knowledge concerning the problem domain is a demanding and tiresome, yet crucial task, executed by knowledge engineers who interview experts. The result of their combined efforts is a knowledge base, the value of which may well exceed the software shell's; and while shells can often be used for different expert systems, permitting leverage of development costs, knowledge bases have to be built over and again for each separate problem area.

One might argue that expert system knowledge as such needs no protection, since the knowledge probably has little value apart from the rest of the system, *i.e.*, without the inference engine to process it. This is only partly true. Even if the knowledge structure is subject to variations according to the needs of the system shell, it may for instance be easier to adapt already formalized knowledge to another shell than to interview an expert from scratch. Infringers could just take the knowledge and re-write the software. In fact, expert system developers themselves often re-write their inference engines in a different computer language once the system has been built and tested. Special artificial intelligence languages such as *Lisp* or *Prolog* ensure comparatively easy system development, but other languages such as "*C*" offer better performance once the system has to be used in practice. It cannot be said off-hand whether a third party who would develop a new shell for existing knowledge would infringe a copyright therein, especially if he were to implement other changes at the same time.

Even more important, expert system developers do not always own exclusive rights to the software shell they use. It is true that software development and knowledge implementation often go hand in hand: one person, or more likely one crew develops the entire system, while only the expert usually steps in from outside. However, the use of standard software shells is increasing. Such shells can be licensed by anybody, including persons who simply intend to obtain the knowledge they need by downloading it from another system. If they do, the copyright in the shell will not protect the knowledge. It will need protection as such. Expert system knowledge therefore needs at least as much protection as the shell. However, in many countries protecting knowledge is more of a problem than protecting software.

IV. Copyright in Knowledge

The Preacher reminds us that "he that increaseth knowledge increaseth sorrow."[8] He may have been referring to other knowledge than such as is commonly used in knowledge based systems, yet many of us will share his concerns, at least when it comes to protecting it. Knowledge is as hard to protect as to obtain. This may especially be true in the case of expert system knowledge. Copyright traditionally focuses on original expression. Such expression is protected while ideas remain free for all. Knowledge can be said to consist of unprotectable ideas; what protection can be obtained must therefore come from the structure.

The kinds of problems we have to face are illustrated by a recent copyright judgment of The Netherlands' Supreme Court. I am aware that from outside, and certainly from the U.S. West Coast, the Netherlands are hard to spot on any globe, and that their law cannot be considered over-important on a world scale. On the other hand, the country is still twice as large as Kuwait, just to mention another nation that didn't quite remain unnoticed lately, and Dutch copyright law is as well developed as any country's. Taking all together, it merits some attention, if only because the issues at stake are the same everywhere. Besides, the very fact that the legal situation may vary from one country to another is an important aspect which deserves separate attention and to which we shall come back later.

V. The Van Dale v. Romme Case

Van Dale Lexicografie publishes the leading general purpose dictionary of the Dutch language. In practice, both the publisher and the dictionary are known as Van Dale. The current (three volume) eleventh edition contains some 230,000 entries or key-words, all of them followed by some explanation or other comments. In preparing it, the editing team profoundly revised the 10th edition's lexicon, removed words that were considered obsolete and added a great number of new ones.

One Mr. Romme, desirous to assist those unhappy puzzle freaks who are bent on solving cryptograms but lack the necessary ingenuity, created a computer program to assist them. The program is meant to provide one who is looking for, say, a nine-letter word meaning, say, what Odysseus forgot when he was staying with Circe, and who merely knows that, say, the fourth and eight letters are a "w" and a "g" with all words meeting these conditions, hopefully including the word "knowledge."[9]

Romme commissioned several typists to copy out all of Van Dale's 230,000 key-words into a program datafile, and even had them add several thousands more from other sources. One supposes the typists must eventually have felt like Hamlet ("What do you read, my Lord?--Words, words, words"), but they did as they were told. The words were not stored alphabetically, as the program required a different order. Still, Van Dale, wishing to reserve to itself any exploitation of its vocabulary, took offence and brought summary proceedings.

---

8 Ecclesiastes I:18.

9 This example of course is fictitious and has no bearing whatsoever on the case.

Both the Utrecht District Court and the Amsterdam Court of Appeal considered the collection of keywords to be copyrighted, and granted an injunction against Romme. According to the Court of Appeal, the vocabulary was original because it was the product of selection by Van Dale's editors who followed their own standards in deciding which words to include and which others to leave out.

However, the Supreme Court of the Netherlands found this reasoning wanting under prevailing copyright principles. In its judgment of 4 January 1991,[10] the Supreme Court begins by repeating the law as it stands: copyright only protects works, i.e., original products which bear the author's personal stamp. It then goes on to say: "A collection of words does not automatically meet this requirement. Such a collection is no more than a quantity of data that is not as such entitled to copyright. This would only be different if the collection should be the result of a selection expressing a personal view of its author." Since it was unclear whether this standard had been applied by the Amsterdam Court of Appeal, the Supreme Court granted the appeal and referred the case to another Appeal Court for further consideration.

VI. The Originality Requirement

Although one must remain aware that the case is not yet over, the judgment demonstrates that data collections may be hard to protect under prevailing copyright rules, not only if they simply contain every relevant item, so that little or no creativity but only a lot of labor went into their being collected, but also if the authors took great pains to collect only those data which they considered sufficiently relevant and leave out the obsolete or unimportant stuff. Such exertions do not necessarily amount to personal creation, at least not in Dutch copyright law.

How about the situation in other countries? As far as I know, all copyright laws require works to be original, but the meaning of this concept varies, at least in detail and at times a great deal more than just that. U.K. copyright law will more easily consider data collections original just because of the skill and labor involved in their compilation. To a lesser extent the same was true for U.S. law until the recent Supreme Court judgment In Re Rural v. Feist.[11] On the other hand, German copyright law is generally understood to be much more severe; in the famous (some would say

---

10 HR 4 January 1991, RvdW 1991, 27 (Van Dale/Romme).

11 S.Ct. 27 March, 1991 In Re Rural Telephone Service Company, Inc. v. Feist Publications, Inc. For earlier cases cf. e.g. NADA v. Business Data, 651 F. Supp. 44 (E.D. Va. 1986) (book listing second hand car prices); Rand McNally v. Fleet Management Systems, 634 F. Supp. 604 (N.D. Ill. 1986) (mileage information in a map) and Koontz v. Jaffarian, 787 F.2d 906 (4th Cir. 1986) (data compilation forming part of a cost estimating system); see also, J.H. Spoor, "Expert Systems and Copyright," in: G. Vandenberghe (ed.), Advanced Topics of Law and Information Technology, Deventer/Boston (Kluwer Law and Taxation) 1989, p. 97 ff.

notorious) Inkassoprogramm-case,[12] the German Supreme Court held that software is only original if more than average programming skills were needed to develop it. The French interpretation of originality is believed to lie somewhere between these extremes, as in fact does the Dutch. Still other approaches may prevail elsewhere.

The very fact that such differences exist is crucial. Literature, music and most works of art are sufficiently original to satisfy the conditions for protection in all countries, even under the more exacting legislations. With data collections this is not the case. Even if such collections are well protected in some countries, the mere potential absence of protection elsewhere forms a serious flaw in the basis for their exploitation as well as their production--at least for such data as are fit for more than local use only.

VII. Special Copyright Regimes for Non-Original Works

In his thesis about copyright in information, Hugenholtz argues after extensive research that "the underpinnings of the prevailing doctrine which denies copyright protection to factual information are weak,"[13] and that "upon critical examination, copyright in compilations of facts is weak and misdirected. Due to the proliferation of information technology, the value of a data compilation is, increasingly, in the data as such, rather than in the way they are presented."[14]

We have to face the fact that present-day copyright is ill-equipped to protecting data collections. As Ginsburg has recently pointed out, "[t]he inhospitability of the personality concept of copyright to fact protection creates uncertain and inconsistent adjudication of claims involving low authorship works."[15] We also have to face the fact that such collections need at least some kind of protection, especially since computer technology made them so much more vulnerable to copying and other forms of misappropriation.[16]

---

12 Supreme Court of the Federal German Republic 9 May 1985, GRUR 1985, p. 1041; cf. also D. Schroeder, Copyright in Computer Programs--Recent Developments in the Federal Republic of Germany, [1986] 3 EIPR p. 89.

13 P.B. Hugenholtz, Auteursrecht op informatie (copyright in information). Deventer (Kluwer) 1989, p. 181. An English translation of this thesis is to be published later this year.

14 Hugenholtz, p. 182.

15 Jane C. Ginsburg, Creation and Commercial Value: Copyright Protection of Works of Information. Columbia Law Review Vol. 90, no. 7 (November, 1990), p. 1865 ff., at p. 1937.

16 Cf., e.g., Ejan Mackaay, Economic Incentives in Markets for Information and Innovation, Harvard Journal of Law & Public Policy, Vol. 13, no. 3 (Summer 1990), p. 867 ff.; Ginsburg, op. cit. and Hugenholtz, op. cit. They all underline the need to protect informational works.

These observations are in no way new. Indeed, more than one century ago, similar considerations led to the introduction into Dutch copyright law of a provision granting a copyright to publishers in all their publications, even those that were devoid of all originality--especially all kinds of lists, directories, etc. Not only did this provision survive, but its importance has even increased in the last few decades, although its field of application remains limited to "writings" only, and moreover only to writings which have been published or are at least intended for publication, these requirements reflecting its origins as a publisher's right. Nevertheless, the prevailing view is that it also applies to data in electronic form, as long as these can be seen as "writings."[17]

Such special regimes exist in several other countries. The details may vary greatly, yet the principle is the same: for certain products, the originality requirement forms a barrier, yet some form of protection is considered appropriate. A well-known example can be found in Swedish law, where directories and other non-original data collections are protected for a limited period by the so-called "catalogue rule." German law has a regime for the protection of minor works, often described as "small change" (kleine Münze). Other examples could no doubt be given.

These special copyright regimes, the details of which vary, still have much in common, as they remain based on some fundamental copyright principles. Essentially, they are obtained without formalities; they give full protection against mere copying, but hardly any against more remote exploitation forms, such as adaptation to a different form; and their duration often is shorter than the term of full copyright.

Many countries have other forms of protection beyond these regimes, based on such concepts as unfair competition and trade secret protection, or provided for in special national legislation. Little can be said about them in general. Differences abound and most of these protection forms can be said to be entirely on their own, at least from an international point of view.

Where does all this lead us in the case of expert system knowledge? Broadly speaking, expert system knowledge is said to consist of facts and rules, most of which represent the generally accepted state-of-the-art in the system's domain. Some of it will more or less reflect one expert's personal views only, although he may be unlikely to acknowledge that those statements

---

17 In the Van Dale v. Romme case this "protection of non-original writings," usually called "protection of writings" for short, did not yet play a role, as both the District Court and the Court of Appeal considered Van Dale's vocabulary to be original. It may yet play a role now that the case has been referred to another Court. There are not yet any examples of its application to data bases, let alone expert system knowledge bases. Examples of its application to other data collections can first of all be found in several Supreme Court decisions with respect to broadcast programming listings, <u>cf.</u>, <u>e.g.</u>, HR 25 June 1965, NJ 1966, 116 ann. HB. For a recent example, <u>cf.</u>, <u>De Toorts v. Oedip</u>, District Court of Amsterdam, 17 May 1989, AMI 1990, p. 51 ann. HCJ, CR 1990, p. 132 ann. Hugenholtz, BIE 1990, 69 ann. CvN, where protection was granted under this regime to a dietary products compendium.

reflect his own personal beliefs rather than hard facts. The presentation of expert system knowledge too is highly formalized, while sequence and structure leave little room for personal choice, if any at all, or so we are told.

So far, the situation is much the same as in Van Dale v. Romme. No doubt those responsible for collecting a dictionary's vocabulary are bound to have their own idiosyncrasies, yet they will do their best to take as neutral a position as possible whenever choices have to be made. In other words, scientists and scholars are as human as the next fellow; still, science is not exactly the ideal setting for turning out products that are stamped with originality.

Taking all together, expert system knowledge may be protected as a copyright work in some countries while falling outside the scope of copyright in others. It may still be protected there by special copyright regimes, if available. Finally, it may enjoy specific protection under all kinds of local laws and doctrines, subject however to all sorts of special formalities or other conditions.

But even if protection may thus be obtained, one highly beneficial condition, if not a vital prerequisite for development and international marketing of new expert systems, is lacking: uniformity of protection. This is an undesirable situation. Endeavors to harmonize knowledge protection, either by way of copyright law or in another form, should therefore be encouraged.

THE ONTOLOGIES OF KNOWLEDGE REPRESENTATIONS: WHY KNOWLEDGE DEFINITIONS SHOULD NOT BE PROTECTED

by

Daniel M. Russell
Systems Sciences Laboratory
Xerox PARC Center
Palo Alto, California
United States of America

Discovering and defining the underlying basis of knowledge representations in AI systems is central to understanding what can and should be protected by legal restriction. To bring out these distinctions, I briefly examine "AI-classical" knowledge in the sense that it exists in expert systems and in various AI programs today. In doing so, we will come to understand what the basic elements of knowledge are, and why a representation's ontology--its key assumptins about the basis of the representation--is definitional and special with respect to sharing knowledge between systems. The implication is that ontologies require legal protections because they have a central role in defining representational systems and because of their potential for supporting knowledge sharing. These protections should not be in the form of patent monopolies, but ought to ensure their status as definitions outside the scope of exclusionary prohibitions.

Representing Knowledge in AI Systems

All languages and representations make certain assumptions about the world they denote. A mechanical engineer designs a device by thinking about objects and their properties, but an artist might draw that device by carefully rendering the negative space around it. Both people are representing the object; each using his own terms and representational techniques appropriate to the task. The basis of the representations, the ontologies, differ fundamentally: the engineer is object-centered, while the artist is manipulating the spaces between objects.

Similarly, if you look in the code or the knowledge base of an AI system, you will see rules, declarations of constants, constraint statements, procedures and so on. A classical expert system works with symbol structures that under certain conditions permit it to draw conclusions from its knowledge base. Not so obvious is the set of assumptions underlying an expert system's representation.

Let's ground this discussion by examining some of the knowledge from a real AI system called Sophie. This is work done a few years ago by Richard Burton and John Seely Brown[1] in their research to build an intelligent tutoring system for basic electronics. Figure 1 illustrates the different kinds of knowledge in Sophie. Sophie knows about things like electronic devices, the component names and some possible ways those parts can fail. Sophie knows how these devices work in the sense that they have particular

---

1 Brown, Burton, DeKleer, "Pedagogical, Natural Language and Knowledge Engineering Techniques in SOPHIE I, II and III" in Intelligent Tutoring Systems, Sleeman & Brown, eds., Academic Press, 1982.

values and they follow particular rules. Sophie can compute with the behavior rules, and even has some troubleshooting rules which allow it to help the student figure out what is going on with the circuit. These classical representation formalisms are used to express knowledge about, in this case, the physics of electronics and electronic device functions. Note here the formalism of that knowledge; that is, the simple syntactic structures used to encode the knowledge. Knowledge is operationally expressed through interpretation with respect to a particular rule context. AI systems do things by interpreting their knowledge structures, that is, going through some process to take those symbolic structures with respect to some semantics about the world. They produce effective computation because the semantics of the representation maintain a correspondence between symbols and the objects and properties being denoted. In classical AI programs, because constant reference is being made to terms within the program that denote external entities, we have a science which delineates between objects in the program, entities, and properties in the world, and tries to keep straight the relationship between the two.

Fig. 1

Device	Possible Faults
resistor	open, shorted, high, low
capacitor	shorted, leaky
diode	open, shorted
zener diode	breakdown voltage high, bv low

Device Parameters

Imin: max allowable reverse current flow, -1 microamp
Imax: max allowable forward current flow, 1 amp
Vmin: min voltage across the diode, -50 volts
Vmax: max voltage across the diode, .8 volts
Ioff: current for diode OFF state, 1 microamp
Ion: current for diode ON state, 2 microamp
Voff: voltage for diode OFF state, .3 volts
Von: voltage for diode ON state, .45 volts

Behavior Rules

GIVEN a new voltage specified as a range, V = [Vl, Vh]
IF Vh <= Voff
THEN propagate the range I = [-infinity, Ion]

IF Vl >= Voff
THEN propagate the range I = [Ion, infinity]

Passive Troubleshooting Rules

IF Ih <= Imin
THEN the diode must be shorted or I must be too low.

IF Il >= Imax
THEN the diode must be shorted or I is too high.

IF Vl >= Vmax
THEN the diode must be open or V is too high.

Figure 1: Knowledge of several kinds in SOPHIE, an electronics debugging intelligent tutoring system.

Those formalisms are the way we write down what we know, and what our programs know, about the world. But there are many assumptions in those knowledge representation structures. There are assumptions about what kind of terms you can use, what are legal forms, what kinds of properties those forms have and the transformations they can undergo as you do computation with them.

All of this is fine, so long as knowledge stays encapsulated within a single system. More and more, however, we are coming to understand that no one application can do everything: multiple sources of knowledge, from a range of developers and programs, must be able to work together. And, as the field of knowledge engineering develops, the ability of systems to interchange knowledge will become increasingly important. Careful consideration of legal protections to enable, rather than discourage, knowledge sharing will become increasingly important both economically (as knowledge must be shared between corportations) and absolutely (as knowledge systems need to expand their knowledge repertory by incorporating knowledge from more and different domains).

Ontologies Allow Knowledge Sharing

We would like to be able to support sharing of knowledge. In order to do so, we have to understand the fundamentals of knowledge representation to see what parts of the knowledge representation can be shared, and what it takes to do so. An example illustrates this best:

In Figure 2 there are three different bodies of knowledge about a pinch roll, a particular device in a photocopier. First the people who do failure analysis of rollers use knowledge about the probability of failure after three months or after six months. Second, the financial people have a cost perspective on this: who is the supplier, information on cost per thousand, cost per hundred thousand, and the difference between these two unit prices. Third, mechanical designers, the people who actually design copiers with these parts, have a different set of representations about pinch rolls dealing with physical properties of the devices, and their subcomponent parts.

Fig. 2

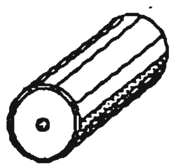

Pinch Roller: X-0025-P-89034-1

Failure Profile

Month	Prob(failure)
1	0.001
2	0.001
3	0.05
4	0.1
5	0.2

o
o
o

Mechanical Properties

(Mass PR1) = 0.2 kg
(Friction PR1) = 2.3
(Modulus PR2) = 9.45
(Thermal-Cond PR1) = 2
o
o
o

Financial/Accounting

Supplier: FoamCore, Inc.
Cost/1000: 229.00
Cost/10,000: 2000.00
o
o
o

Figure 2: Three different kinds of knowledge about a single part. Each is useful to someone (or some system), but in what ways can they be shared?

We would like to be able, even within a single corporation, to share knowledge between these three different sub-organizations. Within Xerox, for example, we have people who have expert systems that know a great deal about mechanical design; we have part systems to figure how to place the rollers inside a copier to get paper from point A to point B without crumpling in the middle. On the other hand, we also have databases and expert systems that know about cost accounting and how most effectively to deliver a product to a customer just-in-time. But there is almost no coordination or integration between these knowledge sources within the corporation.

I mention all this because each database, each source of knowledge, has embedded within it a different set of assumptions underlying the formalism; the representation's view of the world (so to speak) is task-centered and built on the assumptions that aid in problem-solving. That set of assumptions forms its ontology. From the Failure view, a part has a probability of failing as a function of time, from the Mechanical Properties perspective a part has inherent properties (mass, friction, etc.) that attach to the data structure representing the part, while from the Financial/Accounting perspective, a part stands in some relationship with its supplier and its unit cost per order. Thus, each representation has a different basis upon which its representations are constructed: time-ordered functions, object-centered properties, and group participation values.[2]

Ontologies, we believe, are the primary mechanism for understanding how to share knowledge between different AI systems. Although there is disagreement within the field over the details, for our purposes an ontology is "a vocabulary of representational terms with agreed-upon definitions for categories, objects and relations that compose to form a representation."

We like to think of the ontology in both human and machine forms, so, in particular, we have terms and definitions, classes, relations, functions, objects and constants. What are some of the commitments that an expert system has to a set of terms and to a set of definitions, and how do those things work together to represent knowledge? So, an ontology is the set of commitments that some make in order to write down something about the world.

Thus, the failure-profile has representational elements about times, parts and failure probabilities, while the mechanical design representation is concerned with properties of individual components and not with time-varying information. The key point is that each representation is specialized with respect to the task it is intended to perform.

This example illustrates an important point--there is not a single ontology: you can have multiple ontologies. In fact, probably the biggest source of problems in sharing knowledge today is that the people down the street have one ontology for representing knowledge about devices while someone at the other end has a different ontology. Even when the task is the same, ontologies sometimes vary. Representations have different commitments at the term level or at the functional level as to how they represent knowledge. In fact, you can think of a spectrum of ontologies, all the way from very specific ontologies examples (about natural language or about devices) through an ontology which talks about very large categories (like time or causation).

How to Share Knowledge Across Representations

Let us return to our physical device example. Our goal is not only to share knowledge within a corporation, but also between corporations. (You can start to see why this is going to cause intellectual property problems.) In Figure 3 we have business W which makes widgets, and business G which makes gadgets. W has three different sources of knowledge because it is a small company, and all it knows about are the mechanical, electrical, and financial properties of its widgets. So, company W knows how to present facts about the mass of its devices, it knows how to represent basically mechanical information, part dimensions, classical kinds of information that it uses to do testing and dimensioning in machining. Business G, on the other hand, is a different company. It makes gadgets which are bracket-like things which go around the outside of rollers and hold them together. Business G has knowledge about how to service gadgets, information on the financial structure of these things, how much they cost per unit, and also much more sophistication company information. This knowledge can drive the numerical control tools to make these things automatically.

Fig. 3

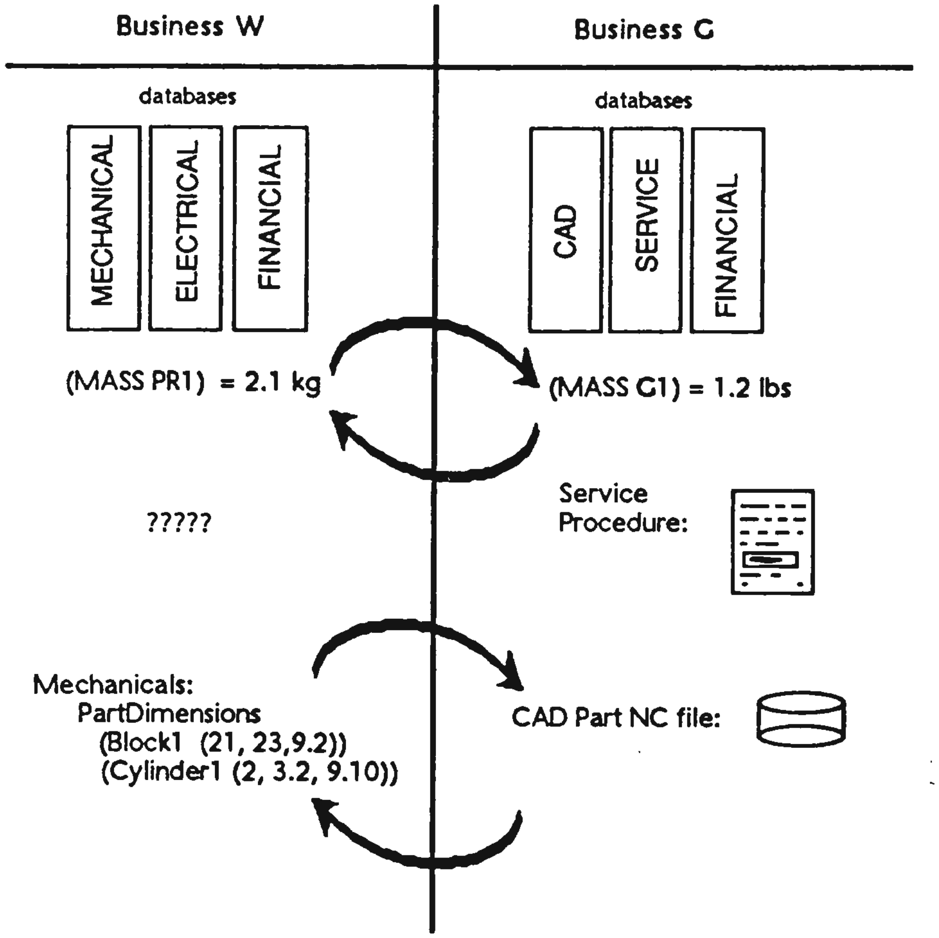

Figure 3: Sharing knowledge between two companies involves simple transformations (e.g., measure conversions from kgs to lbs), and complex conversions (e.g., from dimensioning information and NC programs). In cases where ontologies between data bases are shared, such as the financial databases, conversions are simplified.

Now, what is needed for these two companies to share knowledge to allow some kind of joint venture? In the simplest case, all we have to do is type conversion between different measures. One group uses the metric system, the other uses measures common in the United States. To convert, multiply by the right coefficient and transform the representation from one form to another.

In more complicated cases, (e.g., the service procedure) business G has a kind of knowledge that business W simply does not have. There is no way to transport that without giving these people the fundamental representations that business G has. And, when we get down to the CAD part and mechanical part dimensions, business G has the knowledge but not the procedures; that is, it has the same kind of representations that W has, but does much more sophisticated operations with it. If we want to translate knowledge from one company to another company, we have to identify where the ontologies overlap and where we have similar kinds of commitments.

Let us start at the very bottom and work our way up. At a data level, we can see things like we have ASCII representation or EBCDIC. At a term level, that is, how we name things, we can think about business G here as having terms which, in some cases, are comparable one to one with what business W has. But, in some cases, you have to identify composition of objects or transformations of those objects from one term to another term. The relationships between the electrical components are described in G but not in W, and similarly for service information. Now, in some mechanical kinds of information we have some overlap. There is some sense of shared ontological commitments there, but not completely--each company's representations are based on physical objects and measures of extent along the three dimensions. That is, both share a "physical device in 3-space" ontology, which is probably the most common one going.

A large part of the financial knowledge can be shared. Although there are occasional exceptions,[3] by and large, financial representations are very often based on a common set of ontological assumptions. By and large a debit is a debit and a credit is a credit; entities have balances of money; and money units move from organization to organization in response to promises, contracts and legal agreements. Compared with knowledge in other parts of the organization, financial knowledge can be more easily exchanged and shared because of the sharing in the base ontology.

At the top level, the CAD system overlaps with mechanical data to the extent that they can share representations; they talk about similar kinds of objects. Here, we see, the terms translate one for one through functional transformations, but here we can do CAD operations on mechanical data only to the extent that we can actually reach in and do something functional with that.

To make all of these ideas concrete, let's examine an instance where knowledge was created and could be shared as a consequence of having a common ontology.

---

3 This is not always true; John Sowa of IBM's TJ Watson Research Center tells a story about two mining corporations that combine knowledge bases. The first mining corporation represented all employees with an ID number, while the second mining corporations represented all their employees and all their mules with ID numbers. Wnen these two corporations merged, they did not know quite what to do with this discrepancy. The solution was to adapt the sex field of the employee record so that an employee was either male, female , or mule. Surely an unanticipated extension of the base ontology about genotype!

Common Ontology Use in IDE

The Instructional Design Environment (IDE)[4] is a system that is essentially a knowledge representation workbench; it is a place to create knowledge representation templates (schemata), instantiate the templates, and re-work the representation. These templates define a representation structure, its symbols, relations and terms. In beginning a new task, an IDE user commonly selects a representation from a library, extending it to meet any new requirements.

Fig. 4

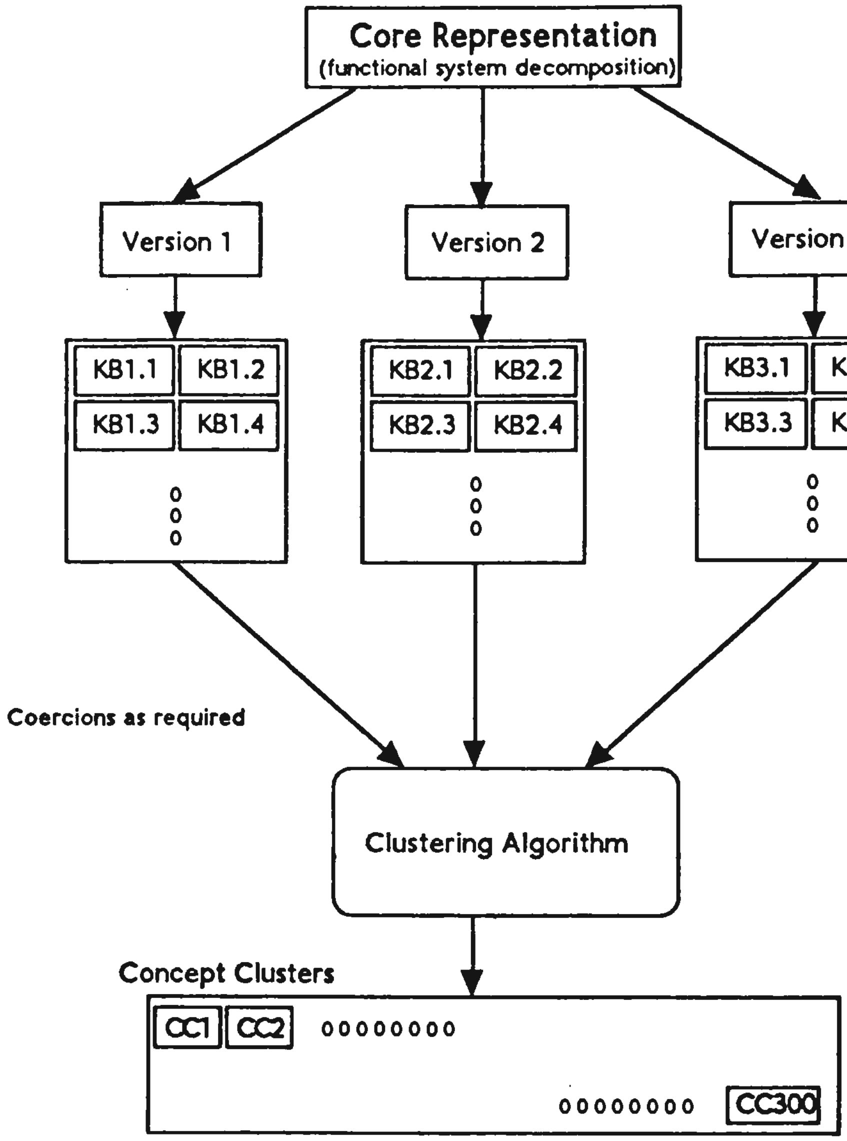

Figure 4: In an actual case of cross-organization knowledge-sharing, having a basic ontology made coordinating and sharing knowledge simpler. Here, a core representation was altered by three organizations, each of which then created seven knowledge bases about laser printers. A clustering algorithm was then used to find correspondences between the twenty-one knowledge bases, and group them into matching conceptual clusters.

---

4 [Russell, Moran, Jordan, 1988] "The Instructional-Design Environment" in Intelligent Tutoring Systems: Lessons Learned, Pstoka, Massey, Mutter eds.; Lawrence Erlbaum Associates, 1988.

Figure 4 illustrates an analysis performed by taking twenty-one different knowledge bases created by three different organizations. Each organization took an element from the representation library, then specialized it for its local use. In each case there was a core representation that was slightly mutated into a locally tailored representation, with each organization's version sharing many parts of the core representation, as well as all of the underlying ontological commitments.

What is interesting about this? This example shows how three different representations can be used by each group to solve its task, and jointly by all to discover concept clusters that no one group understood individually. Because they all came from a common ontological basis, the problems of translation from one form to another were simplified.

Fig. 5

Generality ↓

Reasoning Methods	procedures and algorithms for performing reasoning tasks (examples: constraint satisfactions, theorem-proving)
Ontologies	representational "world views"; commitments between representational elements and world properties / objects / values
Vocabulary	terms and symbols used to denote world properties and objects, often domain specific (examples: the constants and variables within an expert system)
Formalisms	systems of symbols, rules, and structures used to create instances of representations (examples: first-order predicate calculus; frame representation languages)

Figure 5: Knowledge is organized, and can be shared at many levels of detail.

Ontologies are Definitions

I want to summarize with the idea that ontology is fundamentally definitional. Because an ontology establishes the basic terms, relations and structures of a class of representations, it captures a partitioning of the universe in a tractable way without being overly specific. There is a spectrum of ontologies at different levels of abstraction, but there are relatively few high level commitments available to use. The world being the way it is, representations of (for example) physical objects tend to share underlying ontological distinctions. It would be disastrous to the representation industry as a whole if the ontological notions of natural-object, integers, or time-ordered relations (all fundamental ontologies) were patentable. Equivalent versions could possibly be found, but they would be significantly different in kind, expressive ability, and organization, as well as drastically limiting our abilities to share knowledge acoss systems.

Thus, I want to give a sense of an ontology as being in a special category. Perhaps, being so definitional, it is a bit like an algorithm or a natural law. An ontology really makes a commitment to a view of the world, and the way that world-view maps the world onto knowledge representations. I believe one should be able to obtain legal protection for representation formalisms and knowledge encoded within those representations, but I worry about ontologies. Who can own a world view, after all?

Protecting knowledge is one thing, but protecting a gestalt is something else. I believe that it is in the interest of the knowledge representing community at large (and to everyone who will rely on knowledge-sharing between systems) to place representation ontologies, like universal truths, in a special unprotectable category.

IA AND IP: INTELLIGENT APPLICATIONS AND INTELLECTUAL PROPERTY

by

Stephen A. Weyer
Manager, Intelligent Applications
Advanced Technology Group
Apple Computer, Inc.
Cupertino, California
United States of America

Abstract

This presentation and paper highlights "Intelligent Applications", distinguishing between the powerful "engines" provided by artificial intelligence technologies and the useful "vehicles" (i.e., applications) needed and guided by users to reach their goals. Although AI systems could be used in an expert, autonomous mode to create new knowledge and inventions, my emphasis will be on how emerging symbolic and adaptive technologies can be embedded in intelligent applications under user control, for example, information access, CAD and user programming. Given this orientation, can current models of software production, distribution and protection be extended to cover these new kinds of applications?

From AI to IA

"AI" has often meant "Always Impossible"--those things that we don't know how to do yet (and once we do, they are no longer regarded as "AI")--or "Advanced Implements"--the kinds of tools used in rapid prototyping for advanced software systems. The popularly perceived goal of Artificial Intelligence is to rival and perhaps replace people. A more scientific goal is to learn about intelligence, about how people think, and how to simulate this. A third more pragmatic goal is to assist or augment human intelligence. These three perspectives are summarized in Figure 1 below.

Fig. 1

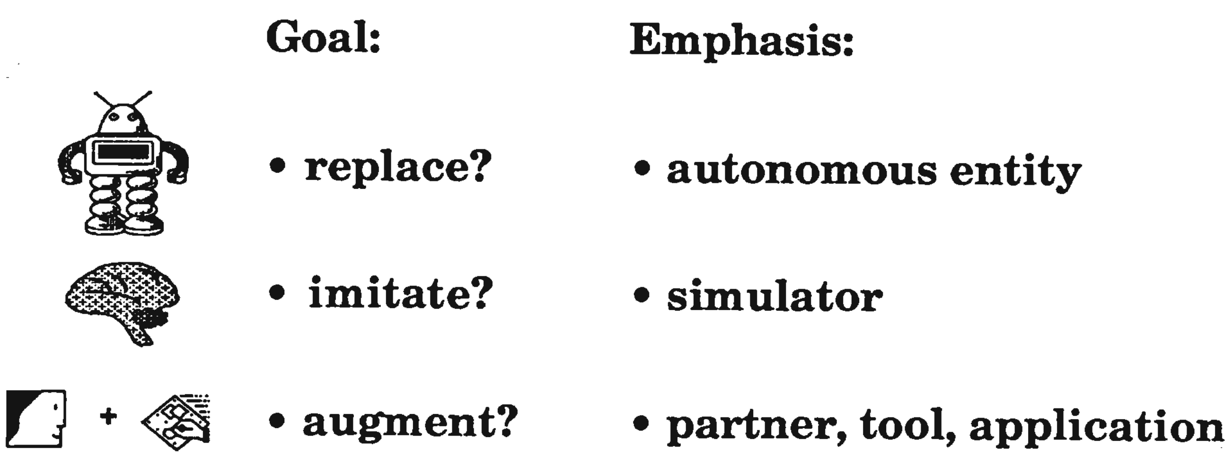

Artificial Intelligence (AI) and Human Intelligence — three perspectives

Rather than emphasize expert or autonomous performance, the goal of "Intelligent Applications" (IA) is to provide useful, everyday tools that support, advise and assist people in problem solving and decision making activities. Rather than worry too much about how to mimic or eventually to replace intelligent human behavior, the focus instead of IA is on hybrid intelligent systems--systems that provide a framework for both human and machine to contribute. Rather than replace people, the goal is to support them, allowing people to delegate routine activities and amplify their abilities. Professor Terry Winograd echoed this IA theme during his presentation: human and computer working together in a "creative partnership", AI as a "new medium". Figure 2 (below) provides a glimpse of how AI technologies might support users, model systems and filter the world of knowledge and data.

Fig. 2

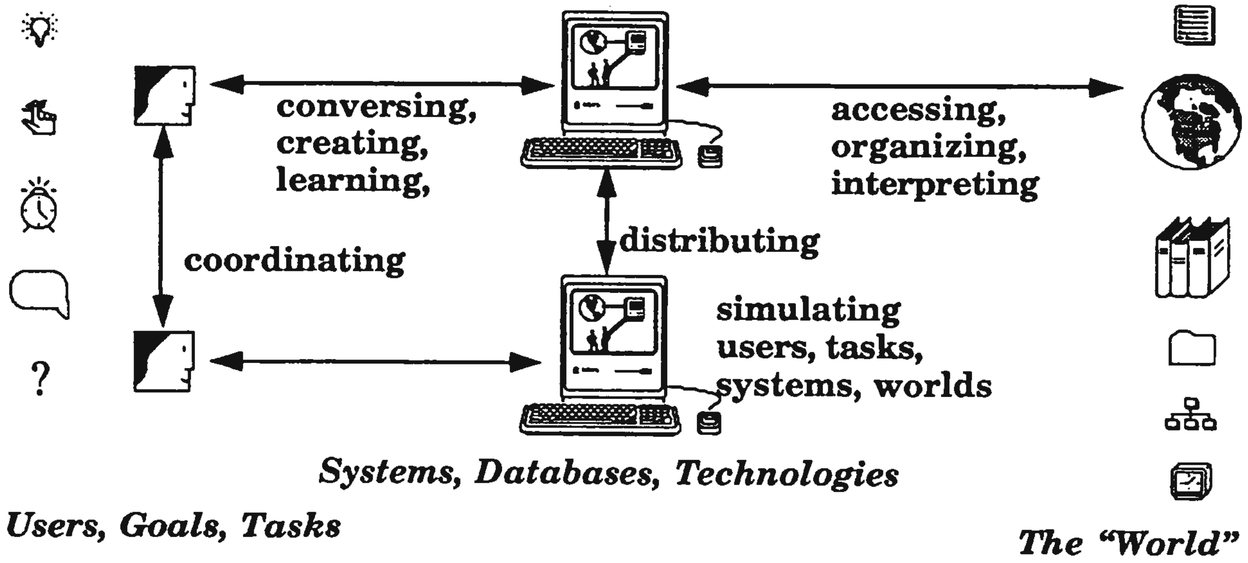

What do Users want from an Intelligent System?
(Asking and Understanding more while Telling it less)

IP and IA/AI

AI technologies provide new and more powerful information engines from which to build new kinds of knowledge vehicles. (We might also ask about the unpredictable, analogous social outcomes of such vehicles, such as information gridlock). Although these systems/vehicles may become increasingly autonomous over time, from a pragmatic and IP perspective, the role of the human user will long be central in the use and creation of knowledge, and the selection and interpretation of results.

We should not rule out the possibility of machine as sole creator or author (and the associated legal issues), but this seems less problematic than some more fundamental issues. For AI to be useful, it must depend on and be integrated with technologies from other areas:

. databases, dictionaries, hypermedia;
. CASE tools, object-oriented programming, software/algorithms;
. hardware accelerators;
. human-machine (user) interfaces.

How clear is our understanding of the technical boundaries and legal protection of these areas? From the other presentations at the symposium, it appears that the law is still rapidly evolving in its interpretations--in particular, Winograd and Davis challenged the prevailing interpretion of software as literary work and as code, and asked for one based on behavior. Even seemingly mundane artifacts such as dictionaries and directories are still controversial.

Assuming for the moment that we have agreed on software interfaces, and databases, do AI technologies differ significantly from these areas? Remember, that once something is understood, it's no longer "AI" but is adopted by the mainstream in computer science and engineering. I do not believe that AI in its current state of capabilities and as incorporated in Intelligent Applications stresses our legal definitions and systems beyond the breaking point.

Given our current model of third-party software developers in the personal computer industry, can ownership of components be retained through licensing, and ownership of extensions or documents created with the tool be claimed by users (who might themselves in turn become 4th-party developers)? Consider a spreadsheet analogy. Users license a basic spreadsheet program from a developer (e.g., Microsoft or Claris). Users (or employers) own the documents that they create using the general program. Application templates for more specific tasks (e.g., financial models for a merger) can be purchased from other developers.

How much protection should such systems have? Here we should attempt to strike a balance that considers at which level standards should be encouraged, how academic freedom can flourish, how users can share informally, and how commercial enterprises can protect their inventions long enough to recoup their investments and leapfrog themselves and others with ever better products.

Figure 3 (see below) is a novice's attempt to sketch a rough map or knowledge space of Intellectual Property. Hopefully, it can serve to suggest questions such as:

. what should the dimensions mean? are they linear?
. How do we compare levels of creativeness or abstractness?
. how do legal protections overlap? where are there holes in protection?
. should we encourage invention and information to be more open?

Fig. 3

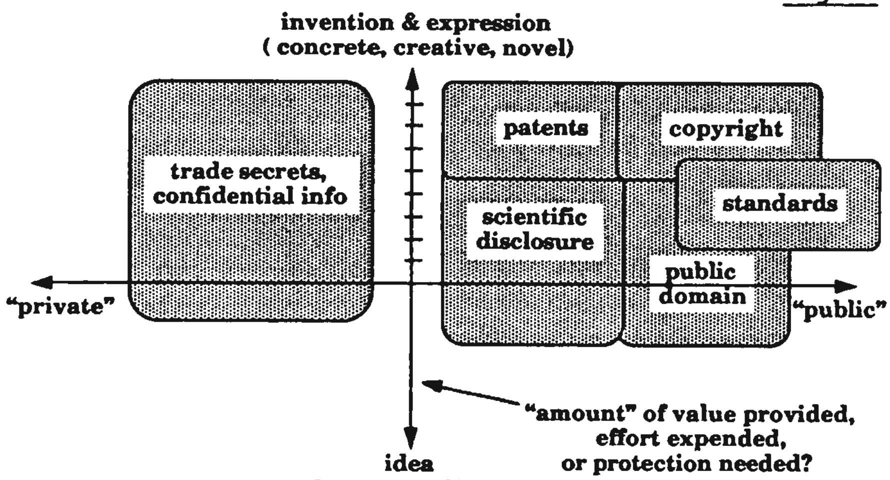

What might an IP knowledge space look like?

Three Example Intelligent Applications and Intellectual Property Implications

1. There are many potential applications of AI technology in the area of intelligent user interfaces: for example, natural language processing, and agents and assistants that model the user (and perhaps have a "personality" of their own). Eager is a system for end-user programming [Cypher91]. A programming assistant notices repetitive patterns in a user's activities and offers to write a script or macro to perform the recognized task on the rest of the user's document. This allows users the option of interacting via both "direct manipulation" (doing it himself/herself) and "delegative management" (asking the machine to do it). Figure 4 (see below) indicates some of Eager's knowledge-based inputs, processing, and results:

- examples (the user's actions) used in training
- the domain (HyperCard™ stacks)
- programming in a scripting language
- the generated program
- knowledge about matching patterns

Which of these can or should be protected?

Fig. 4

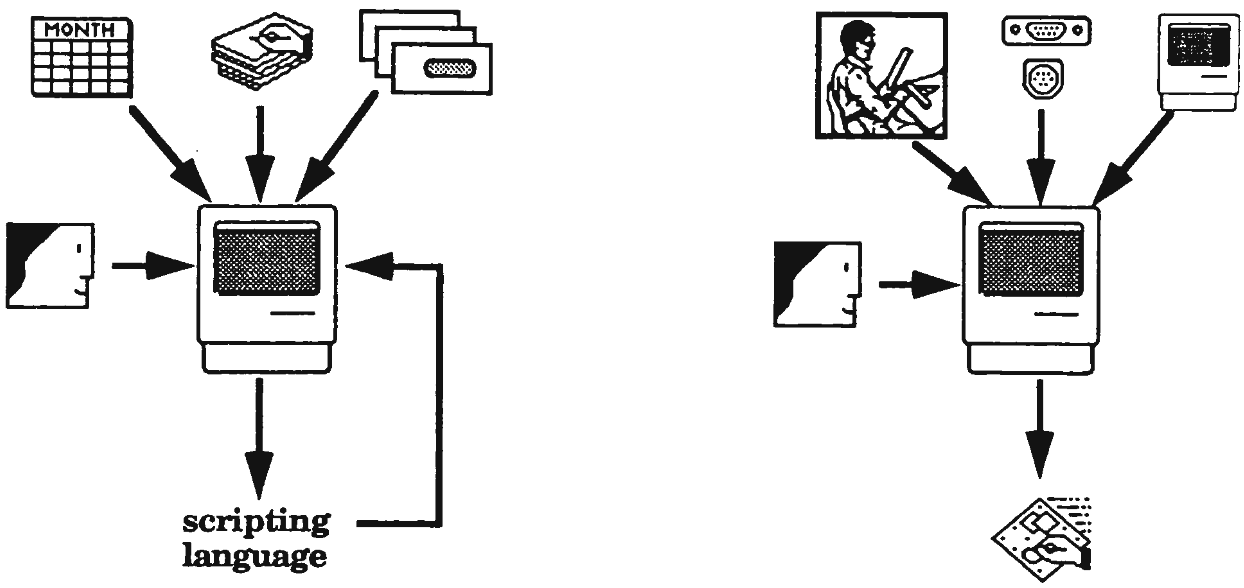

Programming Assistant & Product Design Assistant

2. Consider systems for designing products, configurations of products, or layouts. One approach might be to automate the layout or configuration based on some high-level specification. Another approach, seen increasingly in personal computer CAD software (e.g., Claris CAD or Ashlar's Vellum), is to embed intelligence in an interactive design tool: if the system generates a partial design, the user can edit it; what the user does is checked for consistency and correctness (e.g., violations of any geometric or electrical constraints) by a design assistant. Concurrent Engineering is a related discipline for checking a design against many different requirements in a

product's life cycle: manufacturability, safety, serviceability, etc. These guidelines can be incorporated in an intelligent assistant. Sources of knowledge, as indicated in Figure 4 include:

- design guidelines
- constraint mechanisms
- designs

Which of these can or should be protected?

3. As information becomes available over networks or on large storage devices such as CD-ROM, how are we to access and make use of it? AI technologies could assist in processing the natural language in documents (e.g., news articles, scientific citations, electronic mail messages) and representing the concepts in people's queries. The ObjectLens system from MIT provides rule-based filters for electronic mail. A commercial product, DowQuest [Weyer89] from Dow Jones, is based on a massively parallel computer from Thinking Machines. A free-form query, or headline, or even an entire article (as indicated in Figure 5 below) can be used to create statistically-based filters to locate and rank closest matches from a large document collection.

Fig. 5

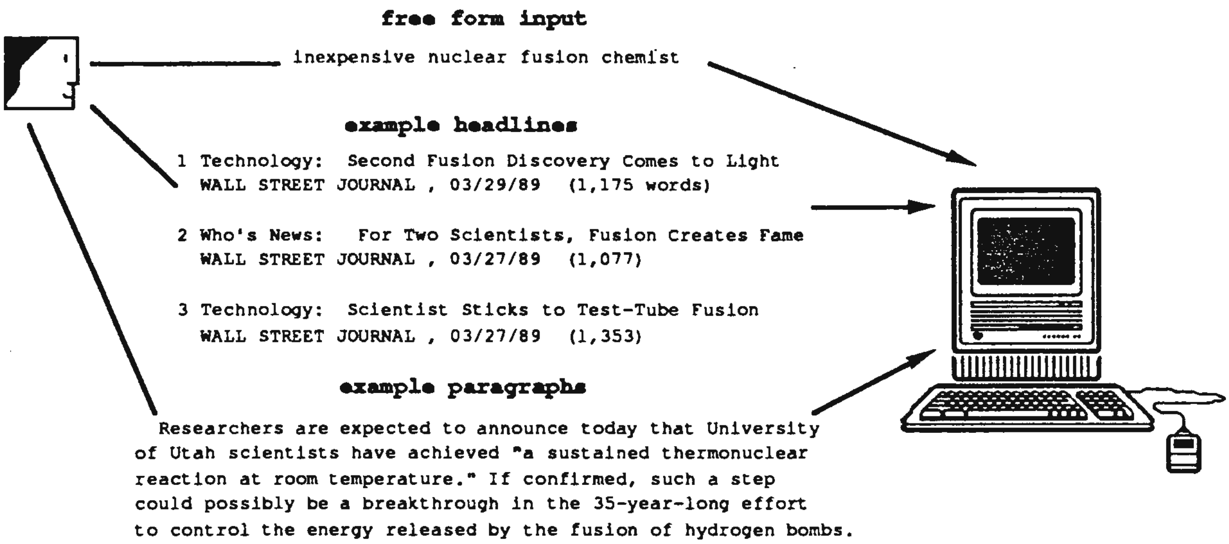

Alternate query styles for DowQuest system.

In addition to the documents themselves, sources of knowledge for such systems could include:

- dictionaries
- grammars
- examples, rules or statistical patterns used in constructing a query

Which of these can or should be protected?

The Future: Creativity or Chaos?

As personal computers and high-capacity networks pervade our planet, and as improvements in AI and other powerful technologies continue to evolve, this could lead to standalone computer-generated works by "AI authors and inventors", but more likely to new tools--Intelligent Applications--to make human inventors and authors more productive and prolific.

On the pessimistic side, this could widen the gap between the technology haves and have-nots. Large numbers of new inventions could overload our administrative and legal systems; tendencies toward overprotection could isolate and reduce the usability of our systems.

On the optimistic side, this outpouring of creativity could lead to a new Renaissance, a flowering in the arts and sciences, improving our environment and quality of life as we enter the 21st century.

Selected References

Cypher, Allen. Eager: Programming Repetitive Tasks by Example. Proceedings of CHI'91. New Orleans, April 28-May 2, 1991; ACM, New York, 1991.

Weyer, Stephen A. Questing for the "Daso": DowQuest and Intelligent Text Retrieval. Online. 13(5): 39-48; September 1989.

INTELLECTUAL PROPERTY AND SOFTWARE: THE ASSUMPTIONS ARE BROKEN[1]

by

Randall Davis
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts
United States of America

1. Introduction

Intellectual property and technology have collided. Again. As has happened repeatedly over several hundred years, the creation and spread of a new technology is forcing rethinking and reassessment of intellectual property law. This time around it is software in general and artificial intelligence in particular that motivate our discussion.

If we are to make progress in the intersection of law and technology, we must understand the essentials of each. We need to understand the fundamental assumptions that underlie intellectual property law and we need to take account of the essential character of software as a technology. Most importantly we need to understand how and where the character of software may conflict with the assumptions that underlie notions of patent and copyright.

In this talk I will argue that one fundamental fact about software distinguishes it from previous technologies: programs are not only text; crucially, they also behave, and it is their behavior that is most often central to these discussions. Because they are both text and behavior, creating programs is simultaneously an act of authorship and an act of invention; programs are both engineered devices and literary works. A central goal of the talk is to make clear this important dual nature of software and to explore some of the important consequences it has for intellectual property law.

I will begin by trying to clear away what I believe to be significant confusion about the issues that face us at this symposium, by defining just a few technical concepts that are central to our undertaking here. With those concepts clearly defined, we can rephrase some of the questions that have arisen, framing them in a way that I believe will expose the crucial issues and help us make progress toward answers.

---

1 An earlier version of a few short segments of this text appeared in the "Musings" column of Chemical Design Automation News, June 1991. The section on digital information draws in part from Pamela Samuelson's article, "Digital Media and the Changing Face of Intellectual Property Law," Rutgers Computer and Technology Law Journal, vol. 16, #2, 1990. The title of the talk is inspired by Allen Newell's article, "The Models are Broken, the Models are Broken," Univ. of Pitt. Law Review, 1986, pp. 1023--1035.

Let me note at the outset, however, than we cannot yet answer the most fundamental question: Should computer programs be copyrighted, patented, both, neither, or something in between? Grappling with that is not on the agenda today, but I believe we can still make considerable progress in setting down the proper foundation.

Instead the agenda is first, to determine what the problem is and isn't. I will claim that the essential problem does not lie in artificial intelligence or any of its variants (like neural nets). The problems we are dealing with arise instead from software in general; nothing essential would be lost if in all of these discussions we were to replace the terms "artificial intelligence" and "neural nets" with "computer program." Almost all the fundamental problems in AI programs arise because they are computer programs, not because they are AI.

Second on the agenda is exploring the character of software, attempting to make clear the assertion that programs are by their fundamental nature simultaneously both literary work and invention.

Third, I wish to explore the consequences of this view. One consequence is the unraveling of some of the abundant confusion that software has generated in the world of intellectual property. A second consequence is the realization that the now infamous litany of structure, sequence, and organization (SSO) is technically incoherent, i.e., it makes no technical sense. While the repair needed to provide coherence is not major, it is important to recognize that the phrase, as it is now understood and widely used, is technically incorrect and only serves to confuse important issues.

Third, I will argue that the fundamental problem is in fact considerably broader than software. It is instead rooted in the digital medium: almost all the hard issues arise in any variety of information in digital form.

Fourth, I will propose an answer to the question raised at the outset of this symposium that gets to the heart of the issue: What do we want to protect against? I will offer a brief answer to that question, one that is surely incomplete, but one that I think will get us well started.

Finally, there remains the question, How shall we proceed? Given some understanding of software and how it strains assumptions underlying intellectual property, we have a significant problem on our hands of reducing or resolving the strains. I have some suggestions about the process, spirit, and mindset in which we attempt to do that.

2. What the Problem Is, and Isn't

Let me begin with the assertion that the problem does not lie in artificial intelligence. I will do that by showing you some work by Harold Cohen, an artist who works at the University of California in San Diego. He is an artist, but his medium is computer software: He writes programs that draw pictures. Figure 1 is of two pictures drawn by a program called Aaron, early in its development.

Fig. 1

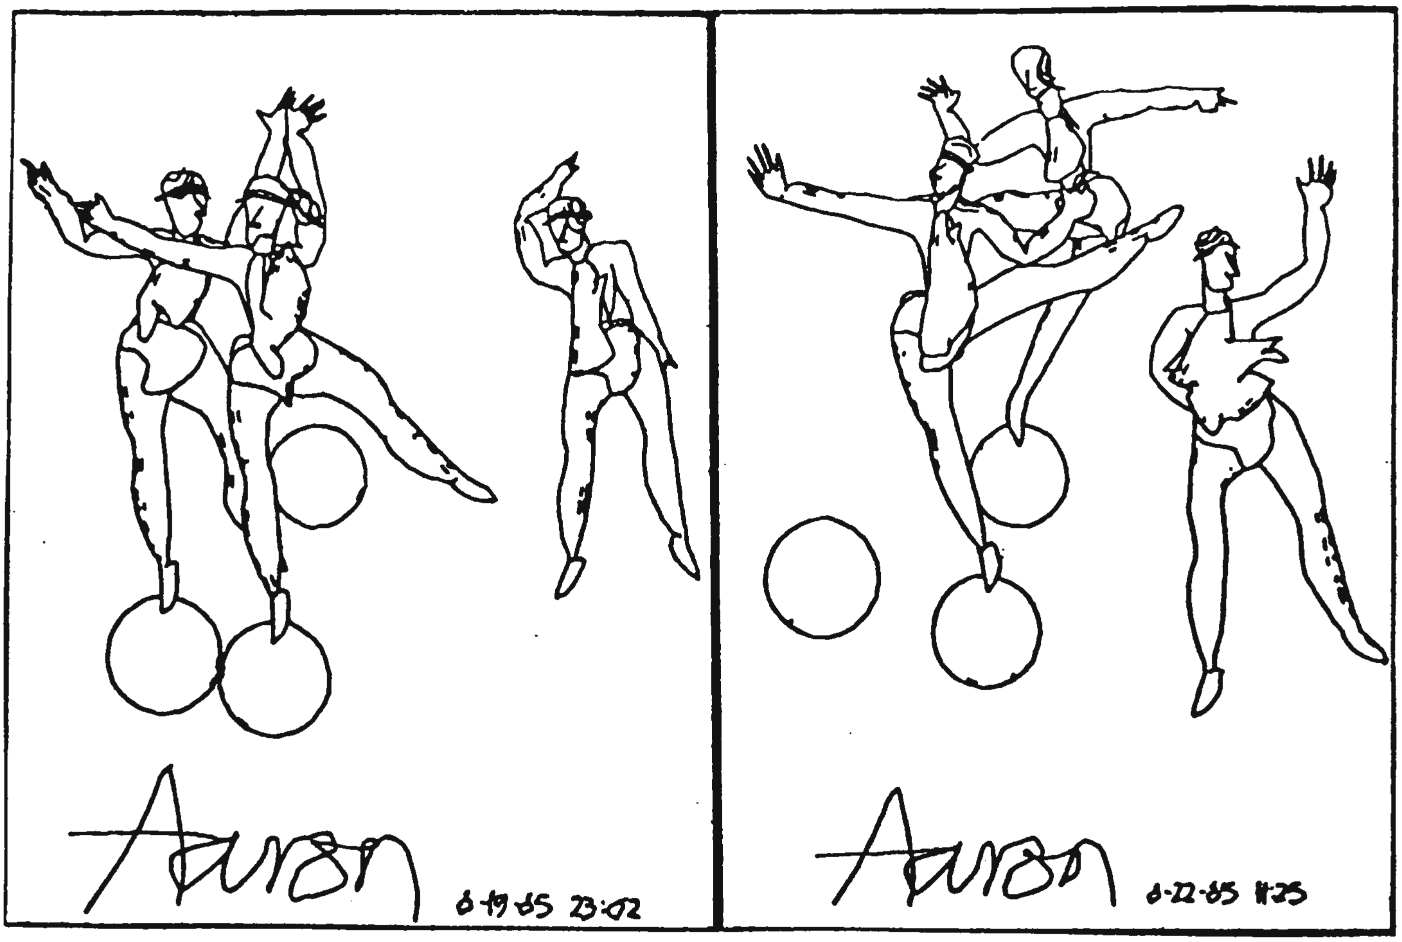

Both of these were drawn at the 1985 AI conference in Los Angeles. Harold had brought the program with him and simply let it run during the conference, generating picture after picture like these and giving them to attendees. The first of them hangs on the wall of my office at MIT; the second I borrowed from a colleague at Stanford.

Note that despite being generated by a machine, neither picture appears "mechanical." The individual human figures appear in graceful and natural stances and the composition--the placement of figures on the page--is aesthetically pleasing. This is particularly impressive in view of Aaron's ability to generate an endless stream of such scenes, with no two precisely alike. Clearly the program is not simply reproducing stored scenes previously drawn. It is instead generating new, different variations each time, and it is doing so based on a set of principles, principles for artistic creation and aesthetics, that enable it to generate work clearly deserving of the term "creative."

This program and its results demonstrate several interesting things. First, it shows that we need not talk about what will happen <u>when</u> or <u>if</u> programs become sufficiently advanced to be creative. They already are. These pictures were generated six years ago; more recent work by Aaron is considerably more sophisticated.[2]

---

2 The cover of these proceedings offers one example, generated by Aaron and colored by Harold Cohen.

Second, computer programs provide particularly rigorous testing grounds for our theories and understanding. A program is a set of explicit instructions that will be executed with exacting precision; hence we discover by simply trying to write a program how much (or how little) we understand a problem, and by running and analyzing the program we can get a detailed and objective determination of which of our ideas provided power in solving the problem. Hence by studying Aaron we can learn something about how creative behavior can be accomplished. They key question about the program is, "What does it know that makes its behavior possible?"

Aaron knows some things about the physical world, such as, humans have two arms, two legs and a head. It knows that human anatomy enables a certain range of motion of the limbs and it uses that knowledge in positioning the figures. It knows that maintaining balance requires putting the center of gravity over the point of support. Notice that all of the figures are well balanced in that way.

The program also knows things about the human world. It knows that there are some joint motions that are anatomically possible but are not normally used (consider the sorts of position brought to mind by the word "contorted"). It knows that humans normally adopt balanced stances, and that we do so by using our arms and legs as counter-balances. That's not the only way to balance: we might also bend from the waist to balance, but we don't normally do that. All of the figures in the picture are in normal, graceful positions, of the sort adopted by dancers, because of what the program knows about the world.

Finally, Aaron also knows things about art and aesthetics. One fundamental principle for composition embodied in Aaron, is "Put it where you can find space." This simple, general principle turns out to be deceptively powerful. To put it to work we need to be more specific: How much space is enough? How far apart should these figures be? What kinds of overlap are allowed?

Aaron's aesthetic of composition specifies answers to these questions as well. With regard to overlap, for instance, Aaron permits partial overlap in figures, but works to avoid obscuring faces. This seems intuitively correct: having one figure obscure, say, a foot, arm, or shoulder of another seems preferable to obscuring the face. If placement of figures leaves no choice and requires obscuring a face, Aaron tries to ensure that the face is not obscured by a foot, because that suggests the unpleasant image of someone being kicked in the head.

3. The Problem Isn't Artificial Intelligence

In Aaron we have a program whose output can realistically be considered to be creative and whose performance derives from a set of explicit and clearly articulated principles for creating aesthetic works.

Hence creative programs exist. What implications follow from this?

For intellectual property I believe the implications are relatively modest. First, human action is still central to this activity: people specify the principles, then write, debug, and run the programs. Determining relative degrees of ownership in the program or its output may be problematic if there are multiple people involved in the construction and running of the program, but that is a familiar issue and one that is independent of AI and software. With human action inevitably at the core of the creative process, the ownership issues seem clear, and more significantly, they remain unimpacted by the nature of Aaron as an AI program.

As I argue in more detail below, the issues that do arise from programs like Aaron come not from the fact that they are AI programs, but from the fact that they are programs. The issues are difficult, intriguing, and are rooted in software in general, rather than in any particular variety of software.

A second significant implication of these programs is their demonstration that we understand something about creativity and aesthetics, an understanding rigorously tested by our ability to build programs that capture some part of it. In particular, we understand it sufficiently well to specify principles for accomplishing it. Those principles happen to be captured in a computer program in this case, but they came from human artists and can be expressed in English, as above. We can as a result use that same sort of understanding to instruct one another in notions of creativity. Indeed we do, as for example in books like A Whack on the Side of the Head: How to Unlock Your Mind for Innovation, R. von Oech, Warner Books, 1988, or The Art of Creative Thinking: A Practical Guide, R. Olsen, Harcourt, 1986. Books like these and programs like Aaron demonstrate that the process is not fundamentally mysterious or impenetrable.

And what in turn is the consequence of that? It is often suggested that mankind is somehow diminished by the ability of machines to perform tasks that seem particularly human. I suggest instead that our ability to create such machines (i.e., such programs) is in fact a significant accomplishment that embellishes rather than diminishes us. The implication is not that creativity and aesthetics are any less remarkable, but that we as scientists and artists are the more remarkable for our ability to understand some parts of those fascinating and subtle phenomena. We not only understand well enough to explain them to each other (as in the books), but understand a few aspects so well that we can "explain" them to a machine, i.e., build programs that capture some of the complexities. While our understanding is surely only a bare beginning, it is still a substantial accomplishment worthy of pride, not an event that diminishes us.

4. The Problem Isn't Neural Nets

There has also been considerable discussion at this Symposium regarding neural nets and the difficulties they may pose. I suggest that for our purposes today, nothing essential is lost if we consider neural nets to be simply a variety of computer program.

A concrete example will facilitate discussion. The simple net in Figure 2 below computes the logic function known as "exclusive-or": its value is 1 if either of its inputs is 1, and is zero if both inputs are 1 or both inputs are 0. There are five component units in the net, interconnected as shown, with weights (of +1 or -1) on the interconnections. The behavior of each unit is the same: if the sum of its inputs is greater than 0, it will emit a 1, otherwise it will emit a 0. The weights on the interconnections influence how the output of a unit propagates to other units: a 1 emitted by unit #2, for instance, will be felt as a -1 input to unit #3 and a +1 input to unit #4. If you try each of the four possible input patterns you can verify that this network computes the exclusive-or of its inputs.

Fig. 2

Exclusive-Or

Inputs		Output
0	0	0
0	1	1
1	0	1
1	1	0

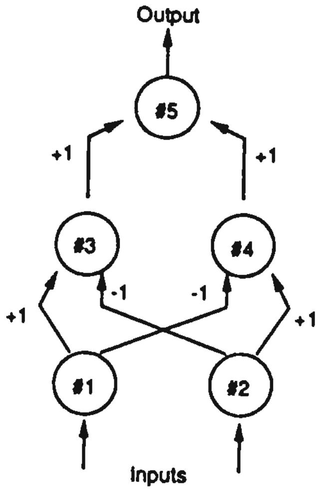

A simple neural network.

More elaborate, real-world applications of networks have been used in areas like credit evaluation, where the inputs are facts about the applicant (e.g., how long have you been at your current job, do you own or rent your home, etc.), and the output is a score indicating the estimated riskiness of the loan.

For our discussion today nothing important is lost if we consider neural nets to be a variety of computer program. These programs are indeed described in a curious fashion, as a set of interconnections and weights, rather than traditional code. They are also created in a curious fashion, typically by being given a set of examples that exemplify what they are to compute, rather than an explicit description of what computation to perform. But neither the creation process nor its results are particularly mysterious. The creation process is a variant of interpolation, fitting a curve to a set of points; more precisely it is a form of non-linear regression. The result of the process--a specific neural net--is a specification for a mathematical function, typically considerably more complex than the "$y = 5x^2 - 3x + 2$" variety we encounter in elementary mathematics, but no more mysterious.

Given a program of this sort, what issues arise for intellectual property? Other speakers have suggested that there may be questions concerning what to protect, since there is no code (at least in traditional, i.e., textual, form), and that difficulties that arise because the weights can be varied to some degree without causing significant change in the performance of the net (hence literal protection of the weights is insufficient).

I suggest that neither of these presents any significant difficulties. To protect a net we need simply protect three things: (i) the pattern of interconnectivity among the units, (ii) the weights on those connections, and (iii) the input and output categories, i.e., the labels that tell us what kind of numbers to put into each input (e.g., in the loan evaluation example, the length of employment), and how to interpret the number(s) that appear at the output (the estimated risk of the loan).

The first two of these can in fact easily be described in the form of a matrix, a table of numbers. The network in Figure 2, for instance, is captured in the table below. To determine the influence that unit #3, for example, has on unit #5, look in row 3, column 5. (Zeroes in the table indicate there is no connection between a pair of units.)[3]

	#1	#2	#3	#4	#5
#1	0	0	+1	-1	0
#2	0	0	-1	+1	0
#3	0	0	0	0	+1
#4	0	0	0	0	+1
#5	0	0	0	0	0

Given the ability of a table of numbers to capture the pattern of interconnections and the weights, all we need to protect is one table and a set of labels. Any real application of a net will have far more units and interconnections than the example above, but the principle is identical: the important behavior of the net is entirely characterizable by a table and a list of categories, i.e., a collection of easily described elements.

The second claimed difficulty arises from the acknowledged fact that, in any sizable network, with dozens of units and perhaps hundreds of connections, the exact weight on each link is not critical. Changes of perhaps 10% seem to have little effect on the performance of the net. In view of that, protecting the exact weights is insufficient. And if small variations on the initial set of weights doesn't degrade performance, how might we protect against someone who copies the original set, varies them randomly by a few percent, then claims independent creation?

One simple answer is already well known: we could easily employ the map-maker's trick of inserting false information into the program. As road maps often carry non-existent streets, so neural nets could be trained to display the initials of the original author when given an obscure or otherwise innocuous set of inputs. Behavior like this from a competing net would give compelling evidence against independent creation.

Hence I suggest that, at least for our purposes today, the problem is not neural nets. We may consider them to be computer programs, programs that are in fact carrying out a rather simple variety of calculation. There are interesting and difficult research problems in neural nets, including understanding how to create nets to carry out specific tasks and determining the limits on the technology. These are interesting unsolved research problems, but they are research issues in artificial intelligence. Here today we are concerned with the result of that process--the computer program that is created--and any problems of intellectual property that arise owe their origins not to the fact that it is a neural net, but to the more general fact that it is a computer program.

---

3 For a good introduction supplying more detail on the general concept of neural nets and the mathematics behind them, see D. Rumelhart et al., Parallel Distributed Processing, MIT Press, 1986, particularly Chapter 2.

5. The Problem Is Software

The difficulties we face, and they are numerous, arise most fundamentally from the character of software as a technology. In this section I consider some of its problematic characteristics and explore how they challenge long-standing assumptions that underlie intellectual property law.

5.1. The Nature of Software: Hardware and Software are Interchangeable

As computer scientists learn early in their education, hardware and software are essentially interchangeable. More precisely, they are what we might call behaviorally interchangeable: any behavior we can accomplish with one we can also accomplish with the other. Any software program can be mimicked, exactly, by a piece of hardware, and the behavior of any piece of digital electronic hardware can conversely be imitated exactly by a program.

For the sake of illustration, consider a software program written to play tic-tac-toe (naughts and crosses). It is easy to write a program that will play a quite credible game, one that will win unless you choose your moves carefully. Given such a program, it would be quite simple to design a machine made from electronic components that played the identical game. That is, the behavior--the choice of moves--of each would be identical.

We could also proceed the other way around: if someone came to us with a machine that played tic-tac-toe (or any other game), we could study the physical design of the machine and create a software program that produced the identical behavior.

There would be <u>some</u> differences, of course: should we decide to change the behavior of either the hardware or software device (perhaps to improve it), we would find the software quite easy to change (typically with a text editor), while changing the hardware version would require the user of soldering irons or wrenches. It is in fact in just this sense that hardware is "hard," i.e., difficult to change, where software is "soft," i.e., malleable.

To show you that this equivalence of hardware and software is more than an abstract principle, I invite you to come to Boston to the Computer Museum where there is machine that will play tic-tac-toe with you. Two curious things about this device are that it is quite large--approximately five feet on a side--and that it is constructed entirely from Tinker Toys and string: every part in it is either a stick or a connector from a Tinker Toy set. Hardware and software are thus not only interchangeable, we need not even use electronic hardware to mimic the behavior of a program.

There are two important consequences of this interchangeability. If hardware and software are behaviorally interchangeable, the choice of which to use in any given circumstance becomes what is termed an "engineering decision." That is, the choice will have no impact on <u>what</u> we can do, only how it will get done. The decision will of course impact things like the cost of the final product, ease of modification, and its speed (hardware is often faster, but not invariably so). So we may well choose between them for reasons of cost, speed, or modifiability, but their fundamental capabilities are identical. Hence nothing fundamental changes as a consequence of switching back and forth between them, and we are free to use whichever one suits our current desires, perhaps even changing from day to day.

The second consequence arises when intellectual property law comes into the picture. While hardware and software are interchangeable in the technical world, notice the enormous difference in the variety of intellectual property protection available depending on which of those we choose. If the device is implemented in software, copyright protection will hold, yet the identical device implemented in hardware will qualify for the far stricter protection offered by the patent regime.

We have here a clash of cultures: a change of no fundamental technical significance leads to enormous variations in the level of intellectual property protection available.

This has lead historically to some interesting anomalies. For example, US Patent #4,135,240, Protection of Data File Contents, offers an improvement in the way a computer allows one user to access the files created by another user.[4] The invention is described as a physical device that works to produce the behavior desired; the diagrams and figures clearly show an electronic device, with logic gates, wires, etc.

But a brief comment in the patent is revealing: "To those skilled in the computer art it is obvious that such an implementation can be expressed either in terms of a computer program (software) or computer circuitry (hardware), the two being functional equivalents of one another .... For some purposes a software embodiment may likely be preferable in practice."

In these two sentences we have, first, a restatement of the principle noted above, and second a completely disingenuous remark concerning the embodiment that "may likely be preferable in practice." In fact no computer system to date has ever implemented this invention in hardware, for a number of good reasons, including the fact that file management is conceptually a part of the operating system, hence this invention belongs in the operating system software, where all the other file management is done. The sham of describing it as a hardware invention was necessary because the patent was filed in 1973, well before Diamond v. Diehr opened the door to software patents in 1981, and well before Apple v. Franklin established in 1982 that copyright protection applied to operating systems. Patent protection for physical devices was of course available, and by minor sleight of hand an invention that is sensibly embodied in software was recast as a hardware device.

The current availability of software patents reduces somewhat the disparity in available levels of protection, but the basic point still stands: for most purposes it is completely immaterial whether a device is implemented in hardware or software, yet intimately associated with each of these is a very different level of protection. Software viewed as a literary work naturally conjures up a copyright approach, while hardware devices are obvious patent material. As a consequence an almost accidental and easily changed technical property of the device has an enormous impact on its intellectual property status.

---

4 For the technical reader: this is the Set UID bit in the Unix operating system.

5.2. The Nature of Software: Programs Also Behave

For some time computer programs have been viewed fundamentally as textual works. In US law, for instance, a program is defined as "a set of statements or instructions ...," bringing to mind of course a written set of statements, and hence a textual work. Indeed a quick glance at any program reinforces this view; they certainly look like textual works.

This is also a seductive view, in part because we already have a body of law--copyright law--well acquainted with textual works. And it almost works: because a program is also a text, many of its text-like aspects are well served by traditional copyright concepts and vocabulary. The confluence of these two produces an almost irresistible tendency to treat software just like any other copyrightable work. But this may be a case of the old saying "When all you have is a hammer, everything looks like a nail." Are programs in fact properly treated as textual works, or are we heading down that path simply because it is familiar and well-worn? There is certainly considerable traffic in this direction, as for example articles like the one by Clapes,[5] whose very title displays the mindset clearly, and whose position is a valiant attempt to see a program as fundamentally a textual work.

A program is indeed a text and many of its text-like aspects can be accommodated by traditional copyright concepts and vocabulary. But the view also causes serious problems, not because it is wrong, but because it is sorely incomplete.

A second, crucially important, thing about a program is its behavior: a program exists in order to get the computer to do something ("bring about a certain result"). A word processing program can for instance, insert text, delete text, copy it, move it, align the margins, etc. Hence a program can also be viewed as a collection of behaviors.

This second aspect of software is crucial: despite the longstanding view of software as a literary work, a program is not only text, it also behaves, and that behavior is an essential aspect of what it means to be a program.

This is particularly interesting because literary works--books--do not behave; behavior is normally associated with machines. Yet programs are textual works created specifically to bring about some variety of behavior. Writing programs is in fact an act of invention.

Put slightly differently,

Software is a "machine" whose medium of construction happens to be "text."

Software is indeed a machine: it is as much a carefully designed artifact as any physical machine and like any physical machine it is designed to do something. Software is indeed a machine: the fundamental interchangeability of software and hardware shows that any program could just as well be a physical machine.

---

5 Clapes, Lynch, and Steinberg, Silicon Epics and Binary Bards: Determining the Proper Scope of Copyright Protection for Computer Programs, 34 UCLA L. Rev. 1493 (1987).

But software is also a unique variety of machine because its construction medium happens to be text.

All this phrasing is quite deliberate, chosen to emphasize assumptions underlying intellectual property law. In the US at least, we have some two hundred years of tradition suggesting a dichotomy between authorship and invention, and the clash of assumption and technology is fundamental.

5.3. Consequence: Creating Software is Both Authorship and Invention

The significance of this for intellectual property is surely clear by now: creating software is, inevitably and inextricably, both authorship and invention. It is also the first technology with this character.

To those of us in the computer science world there is nothing remarkable about this; it is in fact our everyday practice. The clash arises with the tradition in intellectual property law that authorship and invention are distinct activities, with markedly different policies attached to each. In the US the distinction has its origins in the Constitution, with its discussion of "authors and inventors" and "their respective Writings and Discoveries" (emphasis added). The dichotomy was given additional support by the US Supreme Court decision in Baker v. Selden, which has traditionally been interpreted to mean that copyright and patent were mutually exclusive.

The distinction was largely successful: it held strong for some 200 years. But software is different, and it breaks the underlying assumption.

One consequence can be seen in the apparent contradiction that, despite 200 years of tradition that they are disjoint means of protection, software is protected by both copyright and patent. I claim the confusion is entirely appropriate. The source of the problem is not a contradiction in the practice of the law, it is instead the underlying, unjustified assumption that the two concepts are entirely disjoint.

One step further back is the underlying assumption that literary works and inventions are distinct categories. This is almost true; it is the fact that programs also behave that causes the assumption to fail.

5.4. Consequence: The Same Program is Protected by Both Copyright and Trade Secret

Software also presents problems because it is possible to protect the same piece of code by both copyright and trade secret. When programs are distributed, only the machine code is provided; the source code is held as a trade secret. This provides the benefits of trade secret protection: no disclosure and in principle no time limit to the protection. Yet should the source code be revealed, it will be found to have copyright protection. Hence unlike a trade secret, which, once revealed, is public domain, the code is still protected.

This curious combination arises from the confluence of a number of factors, including the fact you can copyright an unpublished work, that software is routinely licensed (not sold), and that the license gives you permission to run the code, but not to read (and decompile) it. While it would make no sense to try to sell a book under the condition that it not be read, it makes fine sense to distribute code in this fashion. This defeats a fundamental underpinning that resides in at least the spirit, if not the letter of the US copyright law: the limited monopoly over reproduction is supposed to "promote the Progress of Science and the useful Arts" at least in part by publication, disseminating the ideas. Yet here we have the monopoly without the dissemination.

Once again, software breaks an important underlying assumption because it behaves. For an ordinary literary work the only way to "use" it is for a person to read it. But code can easily be used without being read by a person. As a consequence a basic mechanism--copyright--is rendered less effective in its intended effect.

5.5. Consequence: The Problem of Software Patents

The patent metaphor similarly runs afoul of essential characteristics of software, particularly in attempts to define patentable subject matter. Historically patents have been denied to laws of nature, "mental steps," formulae, and scientific truths, partly on the grounds that such things are ideas, that ideas cannot be owned (even temporarily), and that patents are supposed to provide incentive to bridge the gap from idea to application. This works when there is in fact a considerable gap between the bare scientific truth and its useful application.

Algorithms are precisely specified methods for accomplishing a task. They are at once objects of considerable mathematical analysis and sophistication (e.g., the study of algorithms for sorting large collections of objects), and they are of immediate practical use. The simplex algorithm of operations research, the Fast Fourier Transform, and the Karmarkar algorithm for the traveling salesman problem, for instance, are all mathematical ideas, but they are mathematical ideas about how to compute something, and as such are immediately useful.

Perhaps the way out then is to deny protection, on the grounds that algorithms are indeed like scientific laws, and on the grounds that incentive for application is not needed when the application is apparent in the idea.

Yet as more and more of the world becomes describable (and described) in computational terms, more and more of what used to be invention in the physical world may get done in software, in the computational world. Computer-aided design systems, for instance, deal not only with shape but are beginning to incorporate descriptions of processes for manufacturing. It may not be long before algorithms become a widely used language for describing (and even inventing) industrial processes, e.g., a process for polishing a rough casting to produce its final surfaces. To deny patentability for such algorithms would be to deny some of the traditional incentive for creating new manufacturing processes, and only because of how it was invented and expressed, not because the process itself was any less non-obvious, or otherwise failed test of patentable subject matter.

The problem is fundamental: the patent concept relies on distinctions and assumptions that do not sit well in the world of software and digital information.

5.6. Consequence: SSO is Technically Incoherent

As we have noted, a program is not only a text; a second, fundamentally important property of a program is its behavior. Each of these gives us a different view of a program: we can look at it as a textual work and we can examine its behavior.

It is crucial to recognize these two views of the program are quite distinct. In particular, they each have their own structure, sequence, and organization.

This is reflected in a useful piece of standard computer science jargon: Programmers talk about both the static structure and the dynamic structure of programs. The static structure is the organization of the program as a textual document (the program-as-text view); the dynamic structure is its behavior, what it does when it actually runs (the program-as-behavior view).

As one illustration of how these two can differ, imagine a word processing program. Two of the things that such a program can do (two of its behaviors) are deleting words and deleting paragraphs. If we were to examine the text of that program, we would likely find in it the sort of thing sketched in Figure 3, with one segment of code that implemented the "delete word" behavior and another that implemented the "delete paragraph" behavior.

Fig. 3

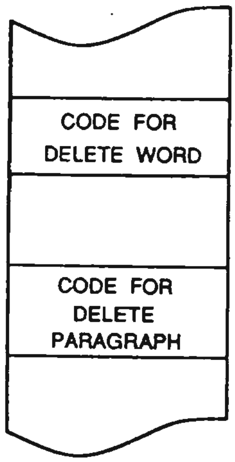

One structure and organization for
the text of a word processor program.

The crucial observation here is that we can quite literally switch the sequence of those two routines in the code (as in Figure 4), but the behavior will be completely unchanged. By that I mean that there is no way that you as a user could determine which of the two programs you were using. If you had been using Version 1 of the program (Figure 3), and without your knowledge we substituted Version 2 (at the right in Figure 4), you would never notice a difference.

Fig. 4

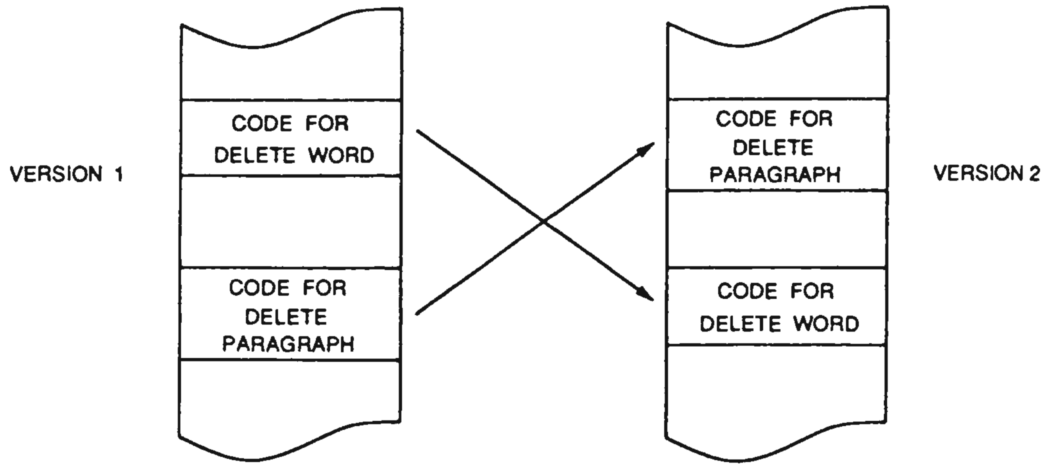

Changing the structure and organization
of the text can leave the behavior unchanged.

This demonstrates that the two aspects of a program--text and behavior--are only loosely connected: we can make significant changes to one without affecting the other. Hence the structure, sequence, and organization of the text of a program is distinct from the structure, sequence, and organization of the behavior of the program.

Another way to illustrate the independence of these two is to imagine using the word processing program long enough that you become quite familiar with the behaviors it offers. Imagine being in the midst of using the program and saying to yourself, "I know very well what this program can do, but I wonder what the code actually looks like?" What might you be able to tell about the text of the program from your experience in using its behavior?

Essentially nothing.

Where, for instance, in the program text is the code for the delete-word command? You can't tell. Does it appear before or after the code for the delete-paragraph command? You can't tell. Hence you may know the structure, sequence, and organization of the behavior of the program quite well, yet be able to draw almost no conclusions at all about the structure, sequence, or organization of the text of the program.

Program text and program behavior are thus only loosely connected, each with its own structure, sequence, and organization.

Note that this is not an artifact of the example we have chosen. It is in fact quite common and arises in part as a consequence of the standard programming practice of dividing programs into independent subroutines.

The independence of program as text (static structure) from program as behavior (dynamic structure) has several significant consequences. The first consequence is that it makes no technical sense to talk simply about the "structure" of a program, because the term is ambiguous and the distinction matters (the dynamic structure can be different from the static structure).

Second, with all due respect for an otherwise carefully argued position, Whelan v. Jaslow[6] runs afoul of this distinction: The (by now infamous) litany of "structure (or sequence and organization)" is flawed by the failure to distinguish between the static and dynamic views of a program. As a consequence the terms are inherently ambiguous, in the important manner noted above.

Indeed, the decision trips coming out of the starting gate, in note 1:

> We use the terms "structure," "sequence," and "organization" interchangeably when referring to computer programs, and we intend them to be synonymous in this opinion.

This is, alas, exactly wrong technically: it is precisely because a program is not only text, but also behaves, that these terms (in all their meanings) are not synonymous. The problem is thus compounded: where the terms were merely ambiguous, they have now been declared equivalent when they are not. The sequence of behaviors can be quite different from sequence of instructions encountered in the program text (recall our ability to change the order in which the subroutines appeared in the program text). The structure of the behavior is quite distinct from the structure of the text (the behavior stayed the same even when the text was reordered).

References to, and hence questions about, "structure, sequence, and organization" as used in Whelan are thus inherently technically defective and as a consequence often unanswerable. Legal use and future decisions may yet provide workable definitions for these terms, or supply better alternatives, but it is important to understand that these currently make no technical sense.

6. What Do We Want To Protect Against?

The question has been raised several times in these discussions, "What is it we want to protect against?" As one attempt to answer it, I suggest that

> A qualifying work of software should be protected against the trivial acquisition of functional equivalence.

It is, first of all, functional equivalence that matters, because the fundamental value in a program lies in its behavior. The question is not whether the code in two programs is identical, or whether the weights in two neural networks are precisely the same; the key is instead behavior. Do the two competing programs do precisely the same thing? If so, then one is a complete substitute for the other and hence it doesn't matter which I use. Thus if you can, in some fashion, take the behavior you have taken the value.

Second, it is the trivial acquisition of functional equivalence that matters. This is important because trivial acquisition means the innovator will have no lead time in which to benefit from the creation, and loss of lead time will then likely lead to loss of incentive.

---

[6] Whelan Associates v. Jaslow Dental Laboratory, 797 F.2d 1222 (3rd Cir. 1986).

Third, note that "acquisition" is purposely vague: how you accomplish it is not particularly significant. The focus here is not the means, it is instead the goal, the valuable goal of equivalent behavior. Acquisition may be by straightforward piracy (i.e., direct copying of either source or object code), it may be via reverse engineering, exhaustive inquiry (in the case of neural nets or data bases), or by any other means now known or yet to be invented. It matters rather little how a functional equivalent is created if the effort involved is trivial.

Hence I suggest that we have the beginning of an answer to the question of what to protect if we focus on functional equivalence as the key property, and recognize that trivial acquisition is the issue, independent of the method by which it is accomplished.

Framing the answer in this fashion has several benefits. It focuses the discussion on the important property of programs (behavior) and the important result of copying (functional equivalence), rather than on the medium of creation (e.g., code, neural nets, or databases). It focuses discussion on the situation we want to avoid (trivial acquisition of equivalence), rather than on the means that can be used to bring about that situation. In focusing on the situation rather than the means, it obtains a degree of technology independence. Reverse engineering, for example, is currently largely manual and hence sufficiently difficult that it does not enable trivial acquisition of any sizable program. Under the suggested principle then, reverse engineering (e.g., "clean room" style re-implementation of a program) would be considered acceptable. But as automated tools for reverse engineering continue to improve, we may someday reach a point where the effort does deserve to be labeled trivial, and at that point that approach may appropriately be judged off-limits. By focusing on the situation, rather than the technology, we have a chance of creating principles that are not soon outdistanced by technological change.

Finally, framing the issue this way also focuses the discussion back on the fundamental issue of incentive. No matter how elegantly written, our principles and laws will be of little use if they lose sight of the goal that brought us to these discussions.

7. The Problem Isn't (Just) Software

Earlier I suggested that the fundamental problems we face in this symposium do not arise out of artificial intelligence, but out of the status of AI systems as programs. In fact the problem is even broader than that.

The problems we face arise most centrally from the digital medium, i.e., any information in digital form. The problems arise here because information in digital form has a number of remarkable properties, properties that are at odds with some of the fundamental assumptions that support enforcement of intellectual property law.

Any information in digital form, is, for instance, orders of magnitude easier and cheaper to replicate. Floppy disks are treated as pieces of paper, yet each of them holds the equivalent of 1500 printed pages and requires only seconds to copy, at close to zero cost.

Digital information is trivial to distribute, worldwide, almost immediately, at very little cost, over an existing infrastructure (i.e., phone lines); we need no additional distribution channels.

Digital copies are indistinguishable from the original; no other medium has that remarkable (and troublesome) property.

Digital information has a breadth of descriptive power we are only beginning to explore: we use it to capture not only software and text, but music, photographs, speech, shape descriptions (e.g., computer-aided design systems), and who knows what else.

Digital information is easily used by multiple people simultaneously: one copy of a file on a large computer can be read by hundreds of users at once and accessed by tens of thousands of people over time at no additional expense.

Digital information is orders of magnitude more compact: a compact disk holds 500 Megabytes, the equivalent of 500,000 printed pages. The information formerly held in 100 cartons of paper has been reduced to the size of something you can easily slip in your pocket.

All of this undermines assumptions built into the enforcement, if not the principle of intellectual property law. What do we do when a wide variety of expensive and complex finished products be duplicated exactly, at effectively zero cost, and distributed almost instantaneously around the planet? When there is no degradation in successive copies, the replication can continue unabated. When copying and distribution is nearly free, there is little to keep the value from slipping out of our hands. When hundreds can read the same copy, there's no need to buy more than one.

Simply put, for other forms of intellectual property, physical law tends to support intellectual property law. The mere weight and volume of the objects copied, the need for distribution channels and physical transportation, etc., all stand in the way of the intellectual property thief. All of that is missing in the case of information in digital form.

The problem isn't (just) software; it's digital information.

8. Designing the Future

What can we do in the face of all this? I suggest that it is crucial to proceed in a non-traditional fashion. Rather than asking "What does the law say?" we need to begin to ask "What should it say?" Put another way, we need to begin treating this as a design problem.

In its simplest terms, a design problem is characterized by a goal, a blank sheet of paper, and some tools. First and foremost is the goal: we must know what we are trying to accomplish. Second is the blank sheet of paper: initially at least, there are no constraints on the solutions we can entertain. Third, the process is driven by the goal, not by the tools. The primary question is not "What do we know how to do?" but "What do we want to accomplish?"

I believe we also need to think of this as a problem of industry design. Any decision about intellectual property law is also a decision about the grounds for competition, and will thereby inevitably shape the future of the software industry. Narrow protection (e.g., protecting only literal source code), for instance, would permit more latitude in creating functionally equivalent programs. This in turn would lead to a marketplace with more software clones and likely shift the basis of competition from functionality to other factors like price, quality of implementation, user support, etc.

Broader protection, by contrast, would likely lead to an increased emphasis on innovative functionality in second comers: being competitive would require functionality sufficiently better that it could win converts from established alternatives.

All this depends on the maturity of the industry as well: early on in most industries there are many small competitors and many new customers. As the industry matures the tendency is toward fewer, larger, well-established companies seeking to lure customers from each other. The best rules for the game thus may differ depending on the stage of industry development.

Hence we are not only attempting to design the future, we are aiming at a moving target. The problem is quite difficult, but I believe there is considerable utility in conceiving of it as a design problem and conceiving of it in terms of industry design.

Consider, for example, the problem of patenting algorithms. Arguments in the legal community have dealt with the somewhat metaphysical problems of the meaning of the term "process," the nature of mental steps, whether algorithms are scientific truths and so forth. We may in fact be better off asking simply, what would happen to the software industry? Would we be better off with or without patents? If we allow software patents, would, as some people claim, software innovation be crippled because of the difficulty of finding all the relevant patents when creating a new system? If we disallow patents do we reduce to any significant degree the incentive for creation of or investment in new ideas? Perhaps the question is in fact better asked in these terms, as one of industrial policy, economics, or psychology, rather than metaphysics.

A similar approach can be take to the question of the scope of copyright protection. Rather invoking the near-mystical principle of idea vs. expression (which even Judge Learned Hand admitted was "inevitably ad hoc"), we might phrase the question pragmatically in terms of lead time: If the time between new releases of a product is typically shorter than the time required to reverse engineer it, the industry will to some extent take care of itself, if we require nothing more than independent creation.

9. Summary

I have argued that artificial intelligence is not currently a significant source of problems for intellectual property law. I have claimed instead that the interesting problems arise from software in general and can only be solved by a clear understanding of the character of software as a technology. One crucial part of that understanding is that programs are not only text, they also behave; software is a machine whose medium of construction happens to be text. Creating programs is as a result simultaneously a work of authorship and a work of invention. In crossing that allegedly unbridgeable barrier software creates significant conceptual difficulties, conflicting as it does with assumptions that have long been a part of the legal system.

The problem is also larger than software alone: digital information of any variety presents intriguing difficult problems for our current intellectual property law framework.

It is in this sense that we need to put aside, even just temporarily, the vast and often valuable experience we have with several hundred years of intellectual property law. For a few moments, at least, we need to examine this problem free of two hundred years of established concepts, categories, and assumptions. In the end it may turn out that some or all of the existing mechanisms of copyright and patent will (once again) prove sufficiently adaptable to gracefully accommodate even this technology. But we need to know that those are the correct tools, not merely that they are the available tools.

And perhaps they will not prove sufficient. Software and digital information generally are unique forms of technology. That uniqueness may have an unavoidable impact on intellectual property law: its clashes with assumptions underlying our current practice may be so profound that some new form of protection will be required. We must have the opportunity to face this question free of preconceptions and the chance to design the future rather than forcing it into molds that were designed in the past.

THE PATENTABILITY OF ARTIFICIAL INTELLIGENCE UNDER U.S. LAW

by

Ronald S. Laurie
Irell & Manella
Menlo Park, California
United States of America

I. INTRODUCTION

Recently, computer programs exhibiting behavioral characteristics which are generally described as "artificial intelligence" have been finding their way into a growing number of business and consumer products. This is in sharp contrast to the situation only a few years ago when the only commercial applications of this software technology were in large computer systems used for complex tasks such as controlling a nuclear power plant, predicting the location of oil deposits, or forecasting next week's weather. Because of the recent emergence of this technology into public view, the intellectual property implications of artificial intelligence have not received much attention in the legal and technical literature.[1]

This paper examines the patentability of artificial intelligence-based, computer-controlled processes and machines by addressing the following questions:

1. What is artificial intelligence and in what respects does it differ from "conventional" (non-intelligent) software?

2. What is the current state of U.S. law concerning the patentability of software-based inventions in general?

3. What new patentability issues are presented by the unique characteristics of technology based on or incorporating artificial intelligence?

II. WHAT IS ARTIFICIAL INTELLIGENCE?

A. The Definitional Problem

Computer scientists do not agree on a rigorous definition of artificial intelligence (commonly referred to as "AI"). Moreover, apart from what AI is, the experts even differ as to whether particular kinds of software systems, e.g., speech or image recognition, are properly characterized as AI or not. This is due at least in part to the fact that the boundaries of AI are continually expanding and contracting in different directions to include new forms of "intelligent" software as they emerge and to exclude others previously considered to be within the domain of AI as they mature. One expert with whom this writer frequently works defines AI, somewhat cynically, as "any software system which is sufficiently sophisticated that it doesn't quite work."

More serious definitions of AI are based on comparing a software system to a human being, either in terms of its operation or its "output," or both. Thus:

> AI is the study of ways in which computers can be made to perform cognitive tasks, at which, at present, people are better.[2]

> Artificial Intelligence programs purport to "capture" or incorporate the logical deductive processes of the human mind.[3]

> The capability to draw inferences is the single quality more than any other [that] constitutes the intelligence of AI software.[4]

> Artificial intelligence involves computer systems that can perform functions which are normally associated with human intelligence, such as learning, adapting, reasoning, self correction and automatic improvement.[5]

An expert's view of AI can depend on his or her scientific vs. engineering bias. Thus, one expert may characterize the primary goal of AI as learning about the nature of intelligence while another may feel that the objective of AI is to build an intelligent machine.[6] Finally, there is disagreement on whether AI processes must mimic or even replicate human thought or whether they need only produce the same kinds of results, i.e., behavior.

> "AI is not really a distinct field of research as much as a family of related areas that share common goals and research techniques. These goals are a desire to understand human intelligence through the building of computer models and/or a desire to build machines that can tackle complex indeterminate problems no matter whether the machines elucidate human intelligence."[7]

The characteristic of complexity applies equally well to conventional programming disciplines, but the characteristic of indeterminacy is what distinguishes AI. The hallmark of AI software is its ability to accommodate uncertainty or ambiguity either in the definition of the problem or in the solution methodology. Whereas conventional computer programming applies a predetermined, and typically repetitive or recursive, algorithmic (step-by-step) solution to a precisely (e.g., mathematically) formulated problem, AI is based on a heuristic (best guess) approach which makes inferences based on stored knowledge (past experience) about relationships between objects having attributes. The relationships, objects and attributes are expressed and manipulated using non-numerical symbols. It is the relationships, and more specifically the heuristic rules defining them, which distinguishes knowledge from data.

B. Categories of AI

Faced with the definitional difficulties referred to above, one turns to the more practical question of what different kinds of computer program-implemented processes are covered by the term. Here, there seems to be general agreement that certain kinds of software qualify as artificial intelligence but there is disagreement as to others. Also, as mentioned above, over time new forms of intelligent software are brought within the ambit of AI while the other more mature forms are banished to the realm of conventional programming. Thus, when "spell checker" software was first developed for use with word processing programs, it was frequently classified as AI, but as the underlying computer science techniques were refined, such software has come to be regarded as more deterministic and therefore as less "intelligent."

Undoubtedly, the category of AI with which the general public is most familiar is the intelligent robot. In their various semi-humanoid forms, these mechanical creatures have been featured in books and film since long before the birth of the electronic digital computer. As is frequently the case, technology eventually caught up with imagination, and science fiction has become scientific fact. Perhaps the next most well-known category of AI, at least to the scientific and engineering communities, is that of expert systems. A group of related software technologies are the "computer-aided" tools such as computer-aided design (CAD), computer-aided engineering (CAE), computer-aided software engineering (CASE), computer-aided manufacturing (CAM) and computer-aided instruction (CAI). AI perception systems include machine vision and speech recognition. Other more embryonic categories of AI are neural networks and natural language generation/processing systems. Finally, there is the very interesting category of computer-generated works which involve little or no human input or control. A brief description of various types of AI software systems is presented in the Technical Appendix at the end of this paper.

Ultimately, the various individual human-like qualities of different categories of AI will be integrated into more comprehensive systems. Thus, the reasoning capabilities of expert systems will be combined with the learning capabilities of neural networks, the perception capabilities of machine vision and speech recognition, and the communication capabilities of natural language systems. The most ambitious efforts to date in this direction are the fifth generation computer projects being undertaken by MCC in the U.S. and ICOT in Japan.

C. Classification of AI Technology by the U.S. Patent and Trademark Office

The U.S. Patent and Trademark Office classification system currently includes only two subclasses specifically devoted to "generic" AI inventions, i.e., those in which the patent claims are not limited to a particular environment, application or end use and those that are not more properly classified elsewhere. Most of such generic AI inventions are found in subclass 513 of class 364 which covers "electrical computers and data processing systems." The official description of subclass 364.513 is:

> Artificial Intelligence (e.g., Speech and Signal Conditioning, Self-Organizing Robot Control); Subject matter . . . wherein the data processing system or calculating computer has the capability to perform functions that are normally associated with intelligence, such as recognition, reasoning, learning, self-improvement, etc. or for solution of problems in this area."[8]

The categories of AI subject matter falling within subclass 513 include generic expert systems, neural networks and complex (computer-assisted) robotic control. The other AI subclass, 513.5, covers intelligent speech processing.

Certain aspects of AI-based inventions are found in a number of classes other than 364. For example, machine vision, including pattern recognition and image processing, is found in class 382 relating to image analysis. As mentioned above, intelligent speech recognition is found in subclass 364.513.5, whereas speech signal analysis not involving data processing is found in class 381 covering audio signal processing generally. The physical construction of robotic inventions is found in class 414 relating to article handling while inventions involving robotic manipulation, e.g., torque control, are found in subclass 318.568.11 relating to programmable electric motor control systems. Inventions involving natural language generation and recognition are found in class 364.419 relating to linguistic applications of data processing. Finally, non-generic, i.e., special-purpose or dedicated AI-based inventions, are assigned to the appropriate class based on the intended environment, application or end use.

While under the current classification system there are only two official subclasses specifically covering AI inventions, the Patent Office has created a computer data base containing a cross-reference collection of patents involving AI which is electronically accessible via two on-line searching systems. This cross-reference art collection is found at subclasses 274-277 of class 364. The AI cross-reference collection is divided into two main sub-categories: software and hardware. Each is further divided into more specific subcategories as follows.

The software subcategory includes expert systems, neural networks, natural language systems, specific programming languages and various AI applications such as robotic control, CAD/CAE design, speech/pattern recognition, medical/electronic diagnosis, weather prediction, instruction, debugging, repair, etc. Expert systems are further subdivided by type (rule-based vs. model-based) and component parts (knowledge base representation, including "fuzzy logic," and inference engine). The hardware subcategories are neural nets, single and plural processors and specific memory techniques (e.g., associative memory).

As part of a comprehensive re-classification of the computer architecture area (which involves some 6,000 to 8,000 patents), the U.S. PTO is reorganizing AI inventions into from sixty to seventy new subclasses. The three major subclasses will be Expert Systems, Neural Networks and Robotics. However, as with current practice these subclasses will be reserved for generic AI inventions, i.e., those as to which the patent claims are not limited to a specific environment, application or end use. Inventions which are so limited will continue to be assigned to other classes, and thus other examining groups. Thus, special purpose (dedicated) expert systems, speech and image processing and CAD, CAE, CASE, and CAM-based inventions, will tend to fall into existing non-AI subclasses determined by the purpose for which the system is being used. This decision reflects the fact that the number of inventions incorporating AI is growing so rapidly that if they were all assigned to a single examining group, the examiners in that group would be inundated with new patent applications. Finally, there will be no distinction between hardware and software. This reflects the artificiality of any such distinction since, at least in theory, anything that can be implemented in software can be implemented in hardware using structured array logic or even random logic. According to the PTO, the reclassification program should be completed in the second half of 1991.

III. PATENTABILITY OF COMPUTER PROGRAM-BASED INVENTIONS GENERALLY UNDER U.S. LAW

A. Constitutional Limits On Patentable Inventions

Legislative authority for the U.S. patent system is found in Article 1, Section 8, Clause 8 of the Constitution which provides that Congress shall have power to enact laws " ... To promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries." The Constitutional clause exhibits a parallelism of language which differentiates between patents and copyrights. Thus, the words "science," "authors" and "writings" form the basis of U.S. copyright law and the words "useful arts," "inventors" and "discoveries" provide the foundation for the patent laws. This clause is both a grant of legislative power and a limitation on that power. Subject to the express and implied limitations contained in the grant, the interpretation of which is for the courts, Congress has the power to select legislative policy which, in its judgment, best realizes the Constitutional goal of promoting progress in science and the useful arts. Also, because the grant is permissive, Congress can enact laws which provide any degree of protection up to the Constitutional limits. As stated by the United States Supreme Court in 1984,

> "As the text of the Constitution makes plain, it is Congress that has been assigned the task of defining the scope of the limited monopoly that should be granted to authors or to inventors in order to give the public appropriate access to their work product. Because this task involves a difficult balance between the interests of authors and inventors in the control and exploitation of their writings and discoveries on the one hand, and society's competing interest in the free flow of ideas, information and commerce on the other hand, our patent and copyright statutes have been amended repeatedly."[9]

B. Statutory Categories of Patentable Subject Matter

Section 101 of the Patent Act of 1952 (Title 35, U.S. Code) provides that any invention (which is defined to include "discovery") in the following categories constitutes patentable subject matter: process, machine, manufacture, composition of matter, or improvement thereof. Section 100 states that the term "process" is synonymous with "art" or "method" and includes a new use of a known process, machine, manufacture, composition of matter, or material. Apart from the substitution of "process" for "art," the language of Section 101 is identical to that in the U.S. patent statute enacted in 1793.

The two statutory subject matter categories which are generally applied to patent claims involving computer programs are process (or method) and machine (or apparatus), although in an appropriate case, a program-related invention might be claimed as an article of manufacture.

In a 1970 decision in one of the first cases to consider the question of whether inventions involving computer programs constitute patentable subject matter, the Court of Customs and Patent Appeals, the predecessor of the Court of Appeals for the Federal Circuit, adopted the phrase "within the technological arts" as the modern equivalent of the Constitutional term "useful arts" and as therefore defining the outer boundaries of patentable subject matter <u>both</u> under the Constitution and under the patent statute,[10]

"... the intention all along has been to convey the same idea and to occupy whatever ground the Constitution permits with respect to the categories of patentable subject matter named in Section 101"[11] (Emphasis added.)

C. Judicially Created Categories of Nonstatutory Subject Matter

Over the years, U.S. courts have identified several types of "discoveries" which are outside the specified statutory categories of patentable subject matter. These include: laws of nature and naturally occurring phenomena (e.g., gravity, electromagnetism, relationship between mass and energy); scientific principles or "truths" (e.g., the Pythagorean Theorem, the relationship of circumference of a circle to its diameter); abstract ideas (e.g., methods of doing business, rules for playing a game); and mental processes (e.g., problem solving techniques).

D. Early CCPA Decisions Involving Computer Programs, 1968-1972

The first cases to consider the question of whether computer programs constitute statutory subject matter were decided by the Court of Customs and Patent Appeals (CCPA) and involved appeals from Patent Office rejections of method claims to software-based inventions on the grounds that the steps of the claimed method described a "purely mental" process. The CCPA first considered the mental steps doctrine in the 1951 case of In re Abrams,[12] although the genesis of the concept is traceable to a 1932 Ninth Circuit Court of Appeals decision.[13] Under this doctrine, if all the steps in a method claim are "purely mental" in character, the claimed invention is nonstatutory.

In a series of cases decided between 1968 and 1972,[14] the CCPA reversed Patent Office Board of Appeals decisions holding that, where process claims were broadly stated such that the individual steps could be performed either mentally or by a programmed digital computer, the claimed subject matter was nonstatutory. In these cases, the Court held that if specific apparatus, such as a computer, was disclosed in the specification, claims which described the process carried out by the computer were statutory, notwithstanding the fact that they could also be performed by the human mind. These decisions effectively ended Patent Office rejections of computer program-related claims on the basis of the mental steps doctrine. However, by this time the U.S. Supreme Court had provided the Patent Office with a new reason for rejecting claims drawn to computer controlled methods or apparatus, namely that they "preempt" a "mathematical algorithm."

E. The Supreme Court Cases: Benson, Flook and Diehr, 1972-1981

In the 1972 case of Gottschalk v. Benson[15] the U.S. Supreme Court considered the patentability of a method of programming a general purpose digital computer to convert signals representing binary-coded decimal numbers to pure binary form. The Court characterized the claimed method as an "algorithm" which it defined as "a procedure for solving a given type of mathematical problem" and cited its prior decision in a 1939 case[16] where it held that "... a scientific truth or the mathematical expression of it is not a patentable invention" (emphasis added). The Court observed that scientific principles, like laws of nature (even though recently discovered), mental processes and abstract intellectual concepts, are "the basic tools of scientific and technological work" and should not be the subject of exclusive rights; whereas their technological application furthers the Constitutional purpose of promoting the progress of useful arts. The Court concluded that claims directed to the particular algorithm for converting binary-coded decimal numbers to binary numbers were not drawn to statutory subject matter--notwithstanding the fact that some of the claims were limited to the conversion of "signals" using a specifically recited digital computer hardware element (a reentrant shift register)--because the mathematical formula had no substantial practical application except in connection with a digital computer. Thus, reasoned the Court, any patent issued on those claims "would wholly preempt the mathematical formula and in practical effect would be a patent on the algorithm itself."

Significantly, the unanimous Supreme Court opinion in Benson made no mention of the CCPA's recent rejection of the mental steps doctrine in favor of its newly formulated "technological arts" standard. (The CCPA had decided that the Benson invention was patentable on the basis that "computers, regardless of the uses to which they are put, are within the technological arts for purposes of /101."[17])

In its second decision on the question of whether patent claims involving computer programs defined statutory subject matter, Parker v. Flook,[18] the Supreme Court in 1978 held that "trivial post-solution activity" will not save a claim which otherwise wholly preempts a mathematical algorithm. In the Flook case, the claim recited a method for updating an alarm limit based on inputs from an industrial process. The claimed method involved the steps of determining the present value of the process variable being monitored, determining a new alarm base according to a specified formula, determining an updated alarm limit using another formula and finally, "adjusting said alarm limit" to the updated value. The Court stated,

> "The notion that post-solution activity, no matter how conventional or obvious in itself, can transform an unpatentable principle into a patentable process exalts form over substance. A competent draftsman could attach some form of post-solution activity to almost any mathematical formula."

The Court also implicitly held that a "field of use" limitation in the claim preamble restricting use of the algorithm to a particular industrial application (catalytic chemical conversion of hydrocarbons) was similarly ineffective in creating statutory subject matter.

The Court in Flook summarized the rationale for its decision in Benson as follows:

> "Reasoning that an algorithm or mathematical formula is like a law of nature, Benson applied the established rule that a law of nature cannot be the subject of a patent." (Emphasis added.)

In 1981, the Supreme Court issued its third (and most recent) pronouncement on the statutory subject matter status of computer program-related patent claims in Diamond v. Diehr and Lutton.[19] In Diehr, the Court held that although mathematical algorithms per-se were nonstatutory,

> ..."a claim drawn to subject matter otherwise statutory does not become nonstatutory simply because it uses a mathematical formula, computer program or digital computer. ...in Parker v. Flook we stated, 'A process is not unpatentable simply because it contains a law of nature or a mathematical algorithm.' It is now commonplace that an application of a law of nature or mathematical formula to a known structure or process may well be deserving of patent protection.

> ...when a claim containing a mathematical formula implements or applies that formula in a structure or process which, when considered as a whole, is performing a function which the patent laws were designed to protect (e.g., transforming or reducing an article to a different state or thing) then the claim satisfies the requirements of Section 101." (Emphasis added.)

The parenthetical reference to "e.g., transforming..." reflected the fact that the claimed method in Diehr was a process for curing rubber which, as a part of the process, used an equation to determine the cure time, i.e. the time the rubber mold should remain closed. Like Benson, Diehr relied heavily on the Supreme Court's 1939 holding in Mackay Radio that, "While a scientific truth or the mathematical expression of it is not a patentable invention, a novel and useful structure created with the aid of knowledge of scientific truth may be." Whereas Benson had used the first half of the Mackay rule to determine the absence of statutory subject matter, Diehr used the second half of the rule to establish the presence of statutory subject matter.

F. Later CCPA Cases: Evolution Of The Freeman-Walter-Abele Two-Step Test, 1978-1982

Based on the Supreme Court's rulings in Benson, Flook and Diehr, the CCPA evolved a two-part test for determining whether a computer program-related patent claim was nonstatutory because it "wholly preempted a mathematical algorithm." The two-step test was first announced in 1978 in In re Freeman[20] and was refined and modified in 1980 by In re Walter[21] and in 1982 by In re Abele and Marshall.[22] This test remains the applicable standard for determining the presence or absence of statutory subject matter in cases involving computer program-related inventions.

In Freeman, the CCPA restated the Supreme Court's holding in Benson as a two step negative patentability test, i.e. a test for the absence of statutory subject matter:

> "First, it must be determined whether the claim directly or indirectly recites an 'algorithm' in the Benson sense of that term, for a claim which fails even to recite an algorithm clearly cannot wholly preempt an algorithm. Second, the claim must be further analyzed to ascertain whether in its entirety it wholly preempts that algorithm."

The Court's reference to reciting an algorithm "in the Benson sense of that term" related to the fact that despite several occurrences in the Benson opinion of the word "algorithm" alone it was clear that the Supreme Court was referring only to mathematical algorithms since any process is an algorithm in the sense that is a step-by-step procedure to arrive at a given result.[23]

In 1980, based on the Supreme Court's holding in Flook two years earlier, the CCPA in Walter modified the second (preemption) step of its test as set forth in Freeman to require a more positive approach to determining what is claimed,

> "If it appears that the mathematical algorithm is implemented in a specific manner to define structural relationships between the physical elements of the claim (in apparatus claims) or to refine or limit claim steps (in process claims), the claim being otherwise statutory, the claim passes muster under Section 101. If, however, the mathematical algorithm

is merely presented and solved by the claimed invention, as was the case in Benson and Flook, and is not applied in any manner to physical elements or process steps, no amount of post-solution activity will render the claim statutory; nor is it saved by a preamble merely reciting the field of use of the mathematical algorithm. (Emphasis added.)

The first sentence in the above-quoted portion of the Walter opinion represented a positive test for the presence of statutory subject matter, although, as discussed below, the CCPA later held in Abele that the Walter positive test was not exhaustive, i.e. it was not the exclusive test for determining the presence of statutory subject matter. The second sentence ("If, however...) restated the negative test in Benson and Freeman, as modified by Flook, and is exhaustive.

In Abele, the CCPA recognized that the Walter formulation of the second (preemption) step of the Freeman test represented two ends of a spectrum:

(a) subject matter which is clearly nonstatutory under the negative/exhaustive test of Benson/Flook, as set forth in Freeman, e.g., the algorithm is not applied in any manner because it is either merely presented and solved or it is claimed in combination with only trivial post-solution activity and/or a preamble field of use limitation; and

(b) subject matter which is clearly statutory under the positive/non-exhaustive test of Walter, e.g., the algorithm is applied in a specific manner to define structural relationships or refine/limit method steps.

The Court in Abele further recognized that there remained a "grey area" of statutory subject matter not covered by either aspect of the Walter formulation,

"... the Walter analysis ... does not limit patentable subject matter only to claims in which structural relationships or process steps are defined, limited or refined by the application of the algorithm.

Rather, Walter should be read as requiring no more than that the algorithm be 'applied in any manner to physical elements or process steps,' provided that its application is circumscribed by more than a field of use limitation or non-essential post-solution activity. Thus, if the claim would be 'otherwise statutory,' albeit inoperative or less useful without the algorithm, the claim likewise presents statutory subject matter when the algorithm is included." (Emphasis added.)

The Court concluded that its broad reading of Walter was consistent with the three Supreme Court decisions, and quoted the language in Diehr (derived from Walter) to the effect that a claim drawn to subject matter which is "otherwise statutory" does not become nonstatutory simply because it contains or uses a mathematical formula, computer program or digital computer.

In effect, the CCPA in Abele defined a positive/exhaustive preemption test ("applied in any manner") as the converse of the negative/exhaustive formulation in Walter ("not applied in any manner"), thereby expanding the

Walter positive/non-exhaustive test ("applied in a specific manner to define ... or refine ..."). In the final sentence quoted above, the Court provided an analytical tool to assist in determining whether the algorithm was "applied in any manner," i.e., remove the algorithm and determine whether the remaining claim elements describe an "otherwise statutory" process or apparatus. The Court then applied its newly formulated preemption test to the invention before it.

The Abele invention was an improved method of generating computerized-axial-tomography (commonly known as "CAT-scan") X-ray pictures, by using a particular image processing algorithm to eliminate unwanted artifacts. Claims 5 and 6 read as follows:

> "5. A method of displaying data in a field comprising the steps of
>
> calculating the difference between the local value of the data at a data point in the field and the average value of the data in a region of the field which surrounds said point for each point in said field, and
>
> displaying the value of said difference as a signed gray scale at a point in a picture which corresponds to said data point.
>
> 6. The method of claim 5 wherein said data is X-ray attenuation data produced in a two-dimensional field by a computed tomography scanner."

The CCPA found claim 5 to be nonstatutory but held that the addition of the limitations in dependent claim 6 created patentable subject matter. The Court felt it was necessary to broaden the Walter formulation in order to reach this result because, "The algorithm, when properly viewed, is merely applied to the attenuation data ... The algorithm does not necessarily refine or limit the earlier steps of production and detection". At this point one might ask, where in claim 6 did the Court find a recitation of "production" and "detection" steps? The answer lies in the Court's observation that the purpose of the two-part analysis is to answer the question, "What did applicants invent?" and in so doing,

> "semantogenic considerations preclude a determination based solely on words appearing in the claims. In the final analysis under Section 101, the claimed invention as a whole must be evaluated for what it is."

Thus, claim 6 required "X-ray attenuation data" and the Court noted that the specification indicated that such data was available only when an X-ray beam is produced by a CAT-scanner, passed through an object and detected upon its exit.

The Court then performed the "otherwise statutory" analysis, "Were we to view the claim absent the algorithm, the production, detection and display steps would still be present and would result in a conventional CAT-scan process."

The Court dismissed the argument that the production and detection steps represented "mere antecedent data gathering" (see Section G1 below) and, mindful of the Flook admonition concerning trivial post-solution activity, pointed out in a footnote that, while the resultant display was an important feature of the invention, the Court's holding did not rest on the "non-triviality" of the display step.

Finally, the Court found that claim 5 was not statutory because the specification provided no greater meaning to "data in a field" than a matrix of numbers "regardless of by what method generated." Thus, the algorithm was neither explicitly nor implicitly (i.e., by reference to the specification) "applied to any certain process." When divorced from the X-ray data production and detection steps implicit in claim 6, the Court found that the display step, by itself, was trivial post-solution activity and, under Flook, could not save claim 5.

In the context of AI-based inventions, it is instructive to compare the basis on which the CCPA found statutory subject matter to be present in Abele with its opposite finding in In re Meyer and Weisman[24] decided a month later. The invention in Meyer was a process and apparatus for testing a "complex system" and analyzing the results of the tests. According to the specification, one application of the invention was in the clinical diagnosis of a neurological abnormality in a patient under examination by a physician, however the claims were not so limited. The diagnostic approach on which the invention was based involved dividing the complex system into a number of elements and associating a factor of function or malfunction with each of the elements. These factors were initialized at the outset and then updated or modified during the course of the process in dependence upon the responses of the system to a series of tests. When the tests were completed, the resultant factors indicated the measure of probability of function or malfunction of the associated elements. The only step in the method claim which described the interaction of the algorithm with the external physical environment was "testing the complex system for a response." It should be apparent that the invention in Meyer was an expert system, although the term is not found in the Court's opinion.

The Court noted that the appellants had conceded that the claims recited a mathematical algorithm representing a mental process that a neurologist should follow. Based on this admission, the Court characterized the invention as "replacing, in part, the thinking process of a neurologist with a computer." The Court based its conclusion that the invention was nonstatutory on its finding that "appellants' independent claims are to a mathematical algorithm representing a mental process that has not been applied to physical elements or process steps, and is therefore not limited to any 'otherwise statutory' process, machine, manufacture, or composition of matter" (emphasis added).

Comparing the results in Abele and Meyer, one may draw the conclusion that in applying the "otherwise statutory" test to the case of antecedent (pre-solution) data gathering steps it may be significant whether the computer-implemented algorithm facilitates, enhances or improves an essentially physical process (e.g., producing X-ray pictures) or an essentially mental process (e.g., subjective diagnosis of an individual by a physician). This issue was revisited by the Court seven years after the Abele and Meyer decisions in In re Grams.

G. CAFC Cases Interpreting The Freeman-Walter-Abele Two-Step Test: In re Grams and In re Iwahashi, 1989

The quartet of 1982 cases involving the patentability of computer program-related inventions,[25] were the last decisions by the CCPA in this area. On October 1, 1982, the CCPA was replaced by the Court of Appeals for the Federal Circuit (CAFC) and it was not until late 1989 that the Court again considered this subject in In re Grams and Lezotte[26] and In re Iwahashi, et al.[27] Grams involved the application of the two step test in the context of method claims combining a mathematical algorithm with one or more antecedent data gathering steps. Iwahashi considered the question of under what conditions method claims and apparatus claims should be considered to be equivalent in determining whether program related inventions are statutory subject matter.

1. The Effect of Antecedent (Pre-solution) Data Gathering Steps: In re Grams

Grams, like Meyer, involved an expert system. The invention in Grams was a process for diagnosing the condition of an individual patient using an algorithm which manipulated data obtained by performing clinical laboratory tests on the patient. The CAFC had to decide whether the claimed diagnostic method was more like the nonstatutory complex system diagnostic algorithm in Meyer or the statutory computer assisted X-ray process in Abele. As in Meyer, the invention in Grams involved a method of testing a complex system to determine whether the system condition was normal or abnormal and if abnormal to determine the cause of the abnormality. Also, as in Meyer, the specification disclosed that the invention was applicable to any complex system, whether electrical, mechanical, chemical, biological etc. but, unlike Meyer, the claims were limited to a clinical diagnostic method for an individual patient. The only independent claim recited as a first step, performing a plurality of clinical laboratory tests on an individual to measure the values of a set of parameters indicative of the individual's condition. The remaining elements of the claim described various processing steps for manipulating the data obtained by the first step. The result of the processing was the identification of a subset of parameters corresponding to a combination of constituents responsible for the abnormal condition. The Court found that these remaining steps represented "a procedure for solving a given kind of mathematical problem," i.e., a mathematical algorithm as defined by Benson. After citing the "otherwise statutory" rule of Abele, the Court in Grams declined to interpret this language literally as declaring patentable any claim that was statutory without the algorithm, referring to the "provided that" qualification in Abele denying statutory effect to field of use limitations or non-essential post-solution activity.

After reviewing a number of its previous decisions, the Court in Grams stated, "Whether Section 101 precludes patentability in every case where the physical step of obtaining data for the algorithm is the only other significant element in mathematical algorithm-containing claims is a question we need not answer."

The CAFC distinguished the Abele case on the basis that there the algorithm served to improve the CAT-scan process and therefore satisfied the Walter requirement of refining or limiting a process step. The Court ignored the language in Abele that "the Walter analysis ... does not limit patentable subject matter only to claims in which structural relationships or process steps are defined, limited or refined by the application of the algorithm" and in fact, in a controversial and somewhat confusing footnote, the Court stated that the Abele rule must be read consistently with Walter "as requiring (to meet the Walter test) not only that the physical steps in the claim (without the algorithm) constitute a statutory process but, also, that the algorithm operates on a claimed physical step" (emphasis added).

The Court concluded that the claimed process was nonstatutory because the step of performing laboratory tests on an individual patient "merely provides the data for the algorithm" (emphasis added). The Court further held that, whether or not the claims required that the method be performed by a programmed computer was irrelevant to the determination of whether the claim defined a Section 101 process.

One possible basis on which to reconcile the Abele and Grams decisions is that in Abele the pre-algorithm step (producing and detecting X-rays), when combined with the post-algorithm step (displaying the X-ray picture), defined a physical process, whereas in Grams (as in Meyer) the algorithm merely replaced the subjective mental processes of a physician. It seems that after lying dormant for almost two decades the "mental steps" doctrine has reemerged sub silentio in the context of AI technology, specifically expert systems.

2. Process Versus Apparatus Claims: In re Iwahashi

Since under U.S. patent law any process can be described in apparatus terms using claim elements in "means plus function" form as permitted by Section 112 of the statute, the question arises as to whether, in the computer program context, the use of process versus apparatus claims produces any difference in result with respect to the presence or absence of statutory subject matter. The traditional rule is that the form of the claim does not control whether the subject matter is statutory because the form of the claim is considered to be a mere exercise in drafting. As stated in Walter,

> "If the functionally-defined disclosed means and their equivalents are so broad that they encompass any and every means for performing the recited functions, the apparatus claim is an attempt to exalt form over substance since the claim is really to the method or series of functions itself. ... In such cases the burden must be placed on the applicant to demonstrate that the claims are truly drawn to specific apparatus distinct from other apparatus capable of performing the identical functions. (Emphasis added.)
>
> If this burden has not been discharged, the apparatus claim will be treated as if it were drawn to the method or process which encompasses all of the claimed 'means'. ... The statutory nature of the claim under Section 101 will then depend of whether the corresponding method is statutory."

In Iwahashi, the claimed invention was an electronic "unit" (i.e. circuit) for providing auto-correlation coefficients used as feature parameters in pattern recognition based on sampled input values. The prior art calculation method was based on a formula involving a multiplication step. The Iwahashi improvement was based on the fact that a close approximation of the correct coefficient values could be obtained without multiplication by obtaining the square of the sum of two of the factors in the equation and calculating the coefficients therefrom according to a stated formula. The advantage of this approach was the elimination of expensive and complicated circuitry associated with implementing the multiplication function. While the overall pattern recognition system may well have involved AI software, the calculation of the autocorrelation coefficients did not, since it was based on a predetermined mathematical relationship.

All of the elements of the Iwahashi claim except one were in means-plus-function form, e.g., means for extracting, means for calculating, means for feeding, etc. The one element which was not in means-plus-function form was a "read-only memory" which contained the squared sum values. The appellants conceded that the claim recited a mathematical algorithm, at least indirectly, however they argued that the second (preemption) part of the two-step test was not met because the mathematical algorithm was implemented in a specific manner to define structural relationships between physical elements of the claim (citing Walter).

The CAFC agreed, finding that the presence of a specifically disclosed and claimed structural element, the read-only memory, resulted in a sufficiently physical implementation of the algorithm. The Patent Office had argued that the recitation of the read-only memory was in practical effect as broad as a means-plus-function recitation (e.g., means for storing...) and that therefore the claim should be treated as being entirely in means-plus-function form. The CAFC disagreed, finding that the read-only memory was a specific apparatus element and further observing that even if it were equivalent to a means-plus-function element, the means-plus-function elements recited in the claim did not encompass "any and every means for performing the functions recited therein." The basis for the latter holding was that, according to the specific language of Section 112, means-plus-function elements "shall be construed to cover the corresponding structure described in the patent specification and equivalents thereof which perform the stated function." The Court observed in a footnote that this rule was applicable not only to determine the scope of a claim in an issued patent but also during Patent Office examination to determine whether a claim reads on the prior art.

The result in Iwahashi seems consistent with the prior decisions of the CCPA. Referring to the previously quoted language from Walter on this subject, the critical inquiry is whether the functionally defined disclosed means and their equivalents are so broad that they encompass any and every means for performing the recited functions. This in turn depends on the level of detail in the specification concerning the preferred implementation. Where the disclosure of specific structure in the specification is very limited, the range of equivalents will be correspondingly broad and the burden will be placed upon the applicant to show that the described apparatus, plus equivalents thereof, do not encompass any and every means for performing the recited function. On the other hand, where the specification discloses specific apparatus, e.g., the read-only memory in Iwahashi, the range of equivalents will be narrower and may not include all apparatus which achieves the stated function.

H. The Patentability Of Computer Program Related Inventions Which Do Not Wholly Preempt A Mathematical Algorithm

It now seems fairly clear that under U.S. law computer-implemented processes and apparatus which do not wholly preempt a mathematical algorithm are not covered by any other of the judicially created exceptions to the statutory categories of patentable subject matter. Thus, in Paine, Webber vs. Merrill, Lynch,[28] a Federal District Court case involving the alleged infringement of a patent covering the concept of using a computer to automatically move funds between several different types of bank accounts to maximize the interest earned, the Court held that the patent was not invalid as claiming a nonstatutory "method of doing business." In reaching this conclusion the Court referred to several of the CCPA decisions discussed above and held that, in the case of patent claims which do not directly or indirectly recite mathematical algorithms (i.e., which do not meet the first step of the Freeman-Walter-Abele test), the only statutory subject matter requirement is that they be "within the technological arts."

Notwithstanding the principle set forth immediately above (and the title of this Section), it is important to remember that computer programs per se are not patentable. Rather, it is only "processes" and "machines" that are implemented or realized using computer programs (without preempting a mathematical algorithm) that are patentable. The program itself, i.e. the series of specific instructions which control the operation of the computer constitute "expression" which is protected by the copyright law rather than the patent law. On the other hand, the actual process performed by the computer (and the virtual machine comprising the computer programmed in a particular manner) is patentable provided, of course, that it is novel and non-obvious. While this concept is useful to keep in mind, it may represent more of a philosophical distinction than a practical one and is the subject of much debate in the U.S. concerning the proper relationship of patent and copyright laws as vehicles for protecting software-based technology.

I. Comment: What's Wrong With Patenting Mathematical Algorithms?

The rule, first announced by the Supreme Court in Benson, denying patentability to inventions which are claimed in such a way as to wholly preempt a mathematical algorithm, is premised upon a fundamental analogy between mathematical algorithms on the one hand and laws of nature and scientific principles on the other. Since the Benson decision, U.S. courts have consistently accepted this premise by assuming that the same policy considerations that preclude patentability for newly discovered scientific principles or natural phenomena--e.g., that material in the public domain should not be removed therefrom--apply equally well to mathematical algorithms. It is submitted, however, that the underlying analogy may not be valid in all cases.

Mathematical algorithms can be used to express basic scientific truths or laws of nature in a particular language, the language of mathematics, using either words or symbols. Thus, the physical relationship between mass and energy may be expressed either in prose or in symbolic form as $E=Mc^2$ and the relationship between the circumference of a circle and its diameter may be expressed as C=/D, or by equivalent words. On the other hand, mathematical algorithms can also be used to express man-made solutions to complex problems, e.g., the optimum allocation of resources within a system consisting of n providers and m users of particular services. The policy considerations which deny patentable subject matter status to the former class of mathematical algorithms do not necessarily apply to the latter.

As discussed above, the Supreme Court derived the Benson rule from its 1939 holding in Mackay Radio that,

> "While a scientific truth or the mathematical expression of it is not patentable, a novel and useful structure created with the aid of knowledge of scientific truth may be." (Emphasis added.)

Clearly, this statement in Mackay was not addressed to all mathematical algorithms (as defined in Benson) but only to those which express a scientific truth. In Benson, the Supreme Court concluded that the effect of patenting the computer-implemented BCD to binary conversion algorithm would be to patent the mathematical formula itself because the algorithm had no substantial practical application except in connection with a digital computer. This could not be done, the Court said, because patenting an algorithm in such a case would be like patenting an idea. This reasoning sounds very much like a variation on the old "mental steps" theme, and in fact, in a footnote in Flook, the Supreme Court, referring to its holding in Benson, stated,

> "... it is not entirely clear why a process claim is any more or less patentable because the specific end use contemplated is the only one for which the algorithm has any practical application."

It has been argued by some computer scientists and legal scholars that computer programs are as much a technological tool in the information age as were steam engines, electric motors and milling machines in the industrial age. The fact that computer programs rely primarily upon the laws of mathematics rather than upon the laws of thermodynamics, electromagnetism or mechanics, should not render them any less patentable, provided the processes and apparatus which they represent are truly new and non-obvious.[29] Under this view of the world, even patent claims which wholly preempt a mathematical algorithm are not necessarily nonstatutory. Rather, the question is whether they define processes or relationships which have always existed but have only been recently discovered or whether, in the words of the Supreme Court in Diamond v. Chakrabarty[30] they are the "product of human ingenuity." The problem of statutory subject matter is then reduced to the question of whether the particular algorithm under consideration is novel in an objective sense.

IV. PATENTABILITY ISSUES PRESENTED BY THE UNIQUE CHARACTERISTICS OF AI TECHNOLOGY

The two bases on which patent claims describing AI-based inventions are most likely to be rejected by the U.S. Patent and Trademark Office are: (1) that the claims do not define statutory subject matter; and (2) that the claims are "indefinite." The sections of the U.S. Patent Statute which would be cited in support of these rejections are Sections 101 and 112, respectively.

A. Statutory Subject Matter

The current state of U.S. law on the patentability of computer program-based inventions generally has been described in the previous section. The implications of those legal rules in the context of AI technology involve the application of the mathematical algorithm preemption rule and the related question as to whether the mental steps doctrine retains any vitality, notwithstanding its express rejection by the CCPA over two decades ago.

With respect to the mathematical algorithm rule, first formulated by the Supreme Court in Benson, AI-based inventions present an interesting question with respect to the first half of the two-step Freeman-Walter-Abele test, i.e., whether the claims directly or indirectly recite a mathematical algorithm. In the few cases involving AI technology (e.g., Meyer and Grams, discussed above), the applicants apparently conceded that the claims described a mathematical algorithm. It is submitted that this was a strategic, as well as a technical, error on their part.

Symbolic processing is an essential characteristic of artificial intelligence:

> "Artificial intelligence is that branch of computer science dealing with symbolic, non-algorithmic methods of problem solving."[31] (Emphasis added.)

An algorithm is a step-by-step procedure, with well-defined starting and ending points, which is guaranteed to reach a solution to a specific problem. Conventional computer architecture readily lends itself to this step-by-step approach, and thus computer programs traditionally have been based on algorithms. Human reasoning processes, however, tend to be non-algorithmic. That is, they consist of more than just following logical step-by-step procedures. As discussed in Section II, AI technology is based on a heuristic approach wherein one estimates the potential of each of the available options which leads towards a particular goal and pursues the most promising option based on past experience, backtracking if necessary, until reaching the "best" solution, i.e., the one with the highest probability of producing the desired result. It should be obvious that this approach is antithetical to algorithmic problem solving methodology. Thus, to the extent that an AI-based invention is claimed in such a way as to highlight its non-algorithmic nature, a strong argument exists that the invention does not satisfy the first half of the Freeman negative test for the absence of statutory subject matter. A supplemental and somewhat less powerful argument is based on another distinction between conventional and AI programming, namely, numeric versus symbolic processing. Conventional programming manipulates data in the form of numbers, whereas AI programming manipulates knowledge in the form of symbols expressing relationships between objects having attributes. Thus, it may be argued that not only is AI software non-algorithmic, it is also non-mathematical.

As discussed above, the mental steps doctrine was unequivocally rejected by the CCPA in several early software cases decided between 1968 and 1971. Under the mental steps doctrine as originally formulated, if a claim described a process that could be performed either by the human mind or by a machine, then its status as statutory subject matter depended on whether the process was "primarily mental" or "primarily physical." It is this aspect of the mental steps doctrine that was rejected by the CCPA. Thus, as long as physical apparatus for performing the process is disclosed, the fact that the claim is phrased broadly enough to read on mental as well as physical implementation does not mean that it fails to recite statutory subject matter. However, as suggested in the previous section, and as illustrated by cases like Meyer and Grams, the kind of thinking which produced the mental steps doctrine may play a part, albeit unconsciously, in leading a court to find that a mathematical algorithm has been preempted. To the extent that this is so, the Iwahashi decision provides some guidance as to avoiding problems based on the quasi-mental character of the claimed invention. Specifically, if the claim is in apparatus (system) form with specific reference to hardware elements which are adequately described in the specification, there should be no statutory subject matter problem, regardless of how closely the machine processes mirror those of a human being.

B. Indefiniteness

Section 112 of the U.S. Patent Act specifies that a patent application "shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as the invention." (Emphasis added.) This is commonly referred to as the requirement of claim definiteness and has some relevance to the patentability of AI-based inventions. The following language from the CCPA's opinion in Musgrave is illustrative:

> "We cannot agree with the Board that these claims (all the steps of which can be carried out by the disclosed apparatus) are directed to nonstatutory processes merely because some or all of the steps therein can also be carried out in or with the aid of the human mind, or because it may be necessary for one performing the processes to think. All that is necessary, in our view, to make a sequence of operational steps a statutory 'process' within 35 USC 101 is that it be in the technological arts so as to be in consonance with the constitutional purpose to promote the progress of 'useful arts.'
>
> "Of course, to obtain a valid patent, the claim must also comply with all the other provisions of the statute, including definiteness under 35 USC 112. A step requiring the exercise of subjective judgment without restriction might be objectionable as rendering a claim indefinite, but this would provide no statutory basis for a rejection under 35 USC 101."[32]

The CCPA's reference to the indefiniteness of claim language which requires "the exercise of subjective judgment without restriction" is particularly pertinent in the context of AI-based inventions. Presumably, the term "subjective judgment" applies only to human thought processes and not to AI. However, as AI techniques continue to evolve and improve in their ability to emulate human thought processes, the line between subjective judgment and computer implemented processes may become, to use an AI term, very "fuzzy" indeed.

NOTES

1 Among the few published articles addressing the intellectual property implications of AI are:

Butler, Timothy L., Can a Computer be an Author? Copyright Aspects of Artificial Intelligence, 4 Hastings Comm/Ent L.J. 707 (1981-1982);

Farr, Evan H., Copyrightability of Computer-Created Works, 15:1 Rutgers Computer and Technology L. J. 63 (1989);

Gable, R. Lewis, Artificial Intelligence: Protection Mechanisms on the New Frontier, 7:5 Computer Law Reporter (July 1988)

Johnson-Laird, Andy, Neural Networks: The Next Intellectual Property Nightmare? 7:3 The Computer Lawyer 7 (March 1990);

Muse, Christopher R., Patented Personality, 4 Santa Clara Computer & High Tech L.J. 285 (1988);

Nycum, Susan H. and Fong, Ivan K., Artificial Intelligence and Certain Resulting Legal Issues, 2:5 The Computer Lawyer 1 (May 1985);

Robinson, Gerald H., Protection of Intellectual Property in Neural Networks, 7:3 The Computer Lawyer 17 (March 1990);

Samuelson, Pamela, Allocating Ownership Rights in Computer-Generated Works, 47:4 U. Pitt. L. Rev. 1185 (Summer 1986);

Spoor, Jaap H., Expert Systems and the Law - An Outline, Copyright (WIPO), pp. 330-336 (October 1990).

2 Encyclopedia of Artificial Intelligence, Stuart C. Shapiro (Ed.), John Wiley & Sons (1990) (Vol. 1) at p. 9 (hereinafter "AI Ency.").

3 Beutel, Richard A., Government Regulation of Diagnostics Software: A Threat to Artificial Intelligence Software Developers, 2:2 The Computer Lawyer 22 (Feb. 1985).

4 AI Ency. (Vol. 1) at page 154 (see Note 2).

5 U.S. Patent and Trademark Office, Patent Classification Definitions - Class 364, pp. 364-36 (December 1989).

6 Understanding Artificial Intelligence, Henry C. Mishnikoff (Ed.), Howard W. Sams & Co,. (1st Ed. 1985) at p. 8.

7 AI Ency. (Vol. 2) at p. 1050.

8 U.S. Patent and Trademark Office, Manual of Classification (June 1990).

9 Sony Corp. v. Universal City Studios, 464 U.S. 417 (U.S. Sup. Ct., 1984).

10 In re Musgrave, 167 USPQ 280 (CCPA 1970).

In a 1980 case involving the patentability of a genetically engineered living microorganism, the United States Supreme Court observed that the legislative history of U.S. patent law indicated that Congress intended that the domain of patentable subject matter includes "anything under the sun that is made by man." Diamond v. Chakrabarty 206 USPQ 193 (U.S. Sup. Ct. 1980).

11 In re Waldbaum, 173 USPQ 430 (CCPA 1972); see also In re Johnston, 183 USPQ 172 (CCPA 1974).

12 89 USPQ 266.

13 Don Lee Inc. v. Walker, 14 USPQ 272 (9th Cir. 1932).

14 In re Prater I, 159 USPQ 583 (CCPA 1968); In re Prater II, 162 USPQ 541 (CCPA 1969); In re Musgrave, see Note 10 supra; In re Benson 169 USPQ 548 (CCPA 1971); and In re Waldbaum, see Note 11 supra.

15 175 USPQ 673 (U.S. Sup. Ct. 1972).

16 Mackay Radio & Telegraph Co. v. Radio Corp., 306 U.S. 86 (U.S. Sup. Ct. 1939).

17 In re Benson, see Note 14 supra.

18 198 USPQ 193 (U.S. Sup. Ct. 1978).

19 209 USPQ 1 (U.S. Sup. Ct. 1981).

20 197 USPQ 464 (CCPA 1978).

21 205 USPQ 397 (CCPA 1980).

22 214 USPQ 682 (CCPA 1982).

23 See In re Pardo and Landau, 214 USPQ 673 (CCPA 1982).

24 215 USPQ 193 (CCPA 1982).

25 In re Abele and Marshall, decided August 5, 1982, see Note 22 supra; In re Pardo and Landau, decided August 5, 1982, see Note 23 supra; In re Meyer and Weisman, decided Sept. 16, 1952, see Note 24 supra; and In re Taner, et al., decided June 16, 1982, 214 USPQ 678.

26 12 USPQ 2d 1824 (CAFC 1989).

27 12 USPQ 2d 1908 (CAFC 1989).

28 218 USPQ 212 (D. Del 1983).

29 See e.g., D. Chisum, The Patentability of Algorithms, 47:4 U. Pitt. L. Rev., 959 (Summer 1986); for a contrary view, see P. Samuelson, Benson Revisited: The Case Against Patent Protection for Algorithms and Other Computer Program-Related Inventions, 39:4 Emory L. J., 1025 (Fall 1990).

30 See Note 10 supra.

31 Buchanan, Bruce G., and Shortless, Edward H., Rule-Based Expert Systems (Addison-Wesley, 1984), page 3.

32 167 USPQ at 289-290.

TECHNICAL APPENDIX

A. Expert Systems

1. Definition

Expert systems include three main components: (1) a knowledge base, (2) an inference engine, and (3) a user interface. The knowledge base contains declarative knowledge, i.e., facts about objects, events, and situations, and procedural knowledge, i.e., information about courses of action, which allow the system to emulate the reasoning processes of human experts in a particular knowledge domain. The inference engine contains the control logic which "navigates" through the knowledge base in response to a particular query input by the user. The user interface typically includes software support in the form of: (a) knowledge editors, which are similar to text editors and word processors and which aid in the creation of the knowledge base, and (b) explanation facilities, which give the user information on how the system arrived at its answer to a particular problem.

2. Basic Construction and Operation

(a) Knowledge Base

The knowledge base contains the declarative and procedural knowledge in the specific area of expertise that the system covers. The two types of knowledge representation most often used in expert systems are rule-based and model-based systems.

In rule-based systems, the declarative knowledge is combined with the procedural knowledge in the form of "IF-THEN" production rules known as "heuristics." (The declarative knowledge is embodied in the IF statement while the procedural knowledge is embodied in the THEN statement.) Heuristics are rules of thumb which simulate the subjective decision making capacity of human experts. These heuristics may incorporate "fuzzy logic" reflecting certain approximations and uncertainties, e.g., IF X is very large, THEN. . . . The difference between fuzzy logic and pure logic is illustrated by the following example: Pure Logic: If A = B and B = C then A = C. Fuzzy Logic: If A is a lot like B and B is a lot like C then A probably equals C.

Model-based systems are often used as a diagnostic tool in determining the cause of a machine malfunction. They contain a model which simulates the structure and proper behavior of the machine. Instead of using "if-then" rules to diagnose a problem, the model-based system compares the machine with the model.

(b) Inference engine

The inference engine contains the logic which decides what rules to apply, and when to apply them in solving the problem. Thus, the inference engine controls the operation of the expert system by: (1) selecting rules to use; (2) accessing and executing the selected rules; and (3) determining when a solution has been found.

(c) User interface

The user interface is the means by which the user communicates queries to the system and the system provides its answers/recommendations to the user. The user interface may also include a knowledge editor and/or an explanation facility (see above).

3. Creation

There are three types of individuals involved in the creation of an expert system: the domain expert, the knowledge engineer and the implementor. Domain experts have significant expertise in a particular area of knowledge but are not necessarily proficient in computer programming or artificial intelligence. The knowledge engineer, on the other hand, is generally an AI expert skilled at developing expert systems.

The knowledge engineer and the domain expert work together to create the knowledge base. The knowledge engineer extracts the appropriate declarative and procedural information from the domain expert and puts it into the knowledge-base in the proper form of representation (e.g., rules in a rule-based systems). The implementor designs the inference engine and integrates it with the knowledge base.

4. Custom v. Standard Systems

The inference engine plus the user interface forms what is called a "shell." Thus, the main expert system component outside of the shell is the knowledge base.

(a) Custom v. Standard Shells

Initially, expert systems were all custom designed since standard expert system software did not exist. Today, however, there are an increasing number of "pre-fabricated" products available. A standard shell can be purchased (licensed) and loaded with expert knowledge. Some software vendors install the particular knowledge desired by a customer while others simply license the shell and the user loads the necessary information. A third option is the so-called "vertical shell." In this case, the knowledge base is pre-loaded with a core of basic information which may be relevant to a number of applications. Additional knowledge relating to a specific narrowly defined area is then loaded by the user.

(b) Custom v. Standard Expert Systems

Expert systems themselves can be categorized as custom or standard. A custom system, which itself may use either a custom or a standard shell, is one which is created to serve the needs of only one customer. A standard system, on the other hand, is one which is designed for and licensed to a number of users. A hybrid form is the expert system designed for a single user but which is useful to others in the same field.

B. Neural Networks

1. Definition

Neural networks come the closest to simulating how the human brain actually functions. The neural nets "learn" how to do a desired task by repeated trial and error testing and self-correction. Whereas, in rule-based expert systems a human expert sets up the heuristics (i.e., the relationships between facts), in a neural net the network is given a set of facts and "guesses" at the correct answer. If the answer is incorrect, the right answer is supplied and the network adjusts its internal decision processes to "correct" the mistake. In effect, the neural network sets up its own rules to relate the input data to the output.

Neural networks can be implemented as software-only, software with some hardware or primarily hardware with software used to communicate with the user.

2. Basic Construction and Operation

(a) Human Neurons

A neural network emulates the action of human neurons in the brain. A human neuron consists of three main parts: the soma, the axon (and its boutons), and the dendrites. The dendrites receive incoming signal pulses from other neurons and transfer these pulses to the soma. The soma "evaluates" all of the signals which it receives from the various dendrites. Some of the signals are positive signals, i.e., they stimulate the soma, while others may be negative, i.e., they act to inhibit stimulation of the soma. To the extent that the soma is stimulated it generates its own electrical signal pulses and transmits them along the axon to the boutons which are connected to either its dendrites or to the somas of other neurons.

(b) Artificial Neurons

The artificial neurons contained in a neural network are designed to mimic the behavior of human neurons. Each artificial neuron consists of parts similar to those found in human neurons. The process starts with an input value. This input may come from an external source (e.g., a keyboard entry) or from another neuron. As each input travels down its dendrite it is assigned a "weight," e.g., +100%, +80% or -90%. These weights are a rough mathematical equivalent of the positive or negative contribution of a human neural electrical pulse to the stimulation of the soma. After all of the inputs reach the soma, they are summed according to their weights, and a resultant net value is produced. This value may require modification in one or more additional summing steps by the soma. The adjusted value is then sent along the axon to other neurons. As it travels down the axon it is assigned a new weight.

A neural network generally comprises three layers of artificial neurons: the input layer, the "hidden" layer and the output layer. The number of input neurons is determined according to the characteristics and number of input data. The number of output neurons is determined the same way. For example, if there are three yes/no questions the answers to which will lead to one of three results then there will be three input neurons (one for each question) and three output neurons (one for each possible result). The number of hidden layer neurons is approximately half the total number of input and output neurons.

The weights are the means which the network relates the input values to the correct output. In a traditional expert system, the knowledge engineer specifies rules and search techniques to correlate input and output. In a neural network, the system itself assigns and adjusts the weights in order to correctly correlate input and output.

3. Creation

Training a neural network is much like teaching humans a foreign language. When the student incorrectly recites a foreign word or phrase the instructor corrects him or her. The student then attempts to correctly recite the phrase and the process continues until the student masters the language.

In the case of a neural network, before training begins the weights from the input layer to the hidden layer and from the hidden layer to the output layer are arbitrarily set. Experience has shown that networks learn better when the weights are initially random and non-zero. During training, a set of inputs with a known output is fed into the network. If the system does not produce the proper output, an error correction is "back-propagated" through the system and all the input/hidden and hidden/output weights are adjusted according to a predetermined formula.

The use of back-propagation to adjust the weights is the most important characteristic of a neural network. The best back-propagation algorithms deliberately introduce randomness and probability into the system. Thus, the human creator of a network may understand the network structure and the algorithms used, but cannot predict the specific internal operations that occur during training.

In actual practice, the training process does not input an entire set of facts before evaluating the output and then back-propagating to correct errors. Rather, the facts are input one at a time and the output from each fact is examined. If an incorrect output is obtained then the weights are adjusted. This results in a shorter learning time for the network.

Several hundreds or thousands of facts may be needed to train a neural net. If a system is badly designed or the training facts are inaccurate then there is a good chance that the network will never learn. Therefore, the programmer's skill in choosing the proper number and type of input and output neurons and the trainer's skill in choosing the proper input factors and data to train the network are critical to the successful operation of the system.

The repeated application of the algorithms to the weights makes it impossible for a human to explain how the weights are assigned. Thus, a neural network learns to represent information in terms that it alone can interpret. Its designer uses intellectual creativity in building the structure and rules by which learning can occur and provides the facts from which to learn, but then essentially loses sight of the specific details of operation and becomes a mere observer. The skill in creating a neural net is in the selection of the data representations (i.e., the number and type of input, hidden and output neurons) and the selection of the training facts.

4. Generic Neural Networks

There are several generic software-only neural networks currently on the market. One such system accepts a user-defined input file describing the number of neurons in the input, hidden and output layer, along with the details of the type of input and output representation required. The user then inputs training facts and the system adjusts the weights, etc.

C. Natural Language Systems

1. Definition

Natural language systems attempt to make communicating with computers comparable to communicating with people. This communication includes typed, printed and displayed language but not spoken language. Oral language processing falls under the AI category of speech recognition and processing.

Natural language processing comprises two distinct areas: understanding and generation. Natural language understanding allows people to express themselves to the computer and natural language generation is the means by which the computer conveys information to the user.

2. Basic Functions

(a) Natural Language Understanding

The problem inherent in natural language understanding is that communications in human languages tends to be ambiguous, imprecise and incomplete. Humans use the context of the communication, their familiarity with the particular type of situation and their expectations about the situation to interpret the communication.

For this reason, accumulated knowledge is indispensable to understanding natural language. It should not be surprising therefore that, as in other AI areas, much of the focus on natural language understanding has been on methods of knowledge representation.

(b) Natural Language Generation

The same problem which applies to natural language understanding also applies to natural language generation, i.e., the system must contain a great deal of knowledge in order to work.

There has not been as much research in natural language generation as there has been in connection with natural language understanding. One reason is that it is easier to figure out what the computer is saying than it is to figure out how to tell things to the computer. However, as expert systems and neural networks become more prevalent, the need to comprehend the system's "thought" process (i.e., the need for the system to generate a comprehensible explanation) will require better language generation capabilities.

D. Computer-Aided Tools and Computer-Generated Works

1. Definition

Computer-aided tools and computer-generated works involve in the creation of a tangible writing (e.g., text, computer program, music, architectural plan, etc.).

2. Basic Operation

(1) Computer-Aided Tools

Computer systems which aid in the creation of a work often use AI techniques associated with expert systems and neural networks. Once again, knowledge representation, the ability to learn and the ability to draw inferences are what transform ordinary computer aids into artificial intelligence systems.

(2) Computer Generated Works

Computer systems which generate works with little or no human input or intervention are, in essence, the ultimate form of computer aided systems. As such, they also use principles of knowledge representation, learning and inference drawing.

3. Applications

Computer-aided tools and computer-generated works are being found in increasingly diverse applications. Computer-aided design (CAD), computer-assisted software engineering (CASE or automatic programming) and intelligent computer aided instruction (ICAI) are examples of systems used to help create works in written form.

Computer-aided design is the process of utilizing the computer to construct drawings or models of objects or systems. Engineering, architecture, construction, electronics, and mechanics are all areas in which CAD systems are currently used. CAD systems do not necessarily include artificial intelligence. The artificial intelligence technology which is most useful in CAD systems is the expert system. For many applications, this use is still in the early stages.

The human input required by computer-aided programming systems ranges from significant to minimal. However, programming systems can be loosely grouped into two categories: intelligent software development tools (AI programs designed to help perform various phases of software development) and automatic programming systems (AI programs which create other programs in response to a human programmer's specifications.)

Currently, software development tools do not, as a matter of course, incorporate artificial intelligence technology. Examples of software development tools include: text editors (which are equivalent to word processors), debuggers (which help programmers test programs to locate and fix errors) and assemblers, compilers and interpreters (which translate higher-level languages into machine-language (binary) code).

Intelligent development tools use AI to enhance the capabilities of conventional development tools. An example of an intelligent development tool is one that incorporates a knowledge base editor, interactive queries and automatic documentation.

The editor contains specific knowledge about programming languages (e.g., proper syntax) and a general knowledge about programming. Based on this knowledge, the editor can detect and correct syntax errors as the programmer inputs the program.

Interactive query is an intelligent enhancement of the debugger. It allows the programmer to query the system concerning the source of an error. Once again, such a system would probably employ expert system technology.

Automatic documentation is an aspect of the query system. It enables the development tool package to automatically generate program documentation, i.e., user manuals and explanations.

Automatic processing is important for two reasons. First, and most obviously, is the usefulness of a system which can accept instructions on the functionality of the desired program and then create it. Second, it is widely believed that automatic programming is itself a critical component of a truly sophisticated artificial intelligent system. An intelligent system must be able to program or reprogram itself in response to a detected mistake, learned principle, outside stimulus, etc.

The ultimate in automatic programming incorporates natural language processing with a variety of expert systems. A natural language user-interface is used to acquire information about the desired process and verify the problem through additional dialogue. This allows the programmer to use ordinary language to describe the desired product instead of having to use formal logic or even examples as the explanation. The various experts are then coordinated to create the program. This coordination is, itself, done by an expert system.

E. Robotics

1. Definition

A robot is a general-purpose machine system, that, like a human, can perform a variety of different tasks under conditions that may not be known a priori.

Robots do not need to look like humans. The main requirement is that they perform tasks requiring flexibility and artificial intelligence. Flexibility refers to their ability to perform a class of different tasks and artificial intelligence refers to the ability of a machine system to perceive "new" conditions, decide what actions should then be performed, and execute these actions.

Flexibility for a given purpose is achieved with a combination of functional components, including: effectors (arms, hands, legs, feet), sensors (which enable the robot to receive feedback from its environment) and auxiliary equipment (tools, jigs, fixtures, tables, conveyors, etc.). Intelligence is achieved through the use of a single top-level computer, called the controller.

It should be noted that non-"intelligent" robots can also be controlled by computer. Therefore, it is important to distinguish between computer-controlled robots and intelligent robots. The primary difference between non-intelligent and intelligent programmable robots is that a non-intelligent robot simply executes preprogrammed motions, while the AI techniques used in programming an intelligent robot allow it to understand its environment and to take appropriate intelligent actions in response to various external situations.

2. Basic Construction and Operation

A robot may be broken down into three main types of components: knowledge representation and reasoning, manipulators and sensors.

(a) Knowledge Representation

Knowledge representation has been discussed in the above sections on Expert Systems and Neural Networks. Another application of these concepts is in the development of robotic programming languages. These languages are designed to expedite programming a robot for a new task or modifying an old one. The robotic programming is accomplished by means of a language processor working in conjunction with the robot controller. The language processor accepts commands (in the robotic programming language) from the user and translates them into commands for the controller. The controller then translates the commands into lower level commands for the proper device (arm, sensor, etc.).

(b) Sensors

Robot sensing is a form of perception, i.e., translation of relevant characteristic or relational object properties into the information required to control the robot in performing a given robot function. The sensing sequence can be divided into two steps: transducing and processing.

Transducing uses hardware to transform the relevant object properties into an electrical signal. These properties can be geometric, mechanical, optical, acoustic, material, electric, magnetic, chemical, etc.

The second step takes the signal and processes it to obtain the needed information. This typically involves two sub-steps. First, preprocessing is used to improve the signal (e.g., filter out noise). This is usually a hardware function. Then the improved signal is interpreted and the necessary information is extracted. This is generally accomplished by software.

Optimum sensing capabilities can not be achieved through a "generic" sensing strategy. Instead, the steps should be performed by the software/hardware combination best suited to the particular sensing task. This combination will change based upon the particular object properties and the desired resulting robotic functions. Therefore, a sensing system should have a variety of tools and a knowledge-based supervisor (i.e., an expert system) to coordinate the sensing process.

(c) Manipulation

Robot manipulation entails the kinematics, motion trajectories, dynamics, and control of a robot arm. For example, kinematics is used in connection with the position and orientation of a robot wrist (i.e. calculations must be made to determine the angles and lengths of the arm sections in order to arrive at the proper wrist location). Motion trajectories determine the arm's path of movement. Dynamics applies the proper forces/torques to the joint motion and Control makes sure that the torque and the motion are synchronized.

PATENT APPENDIX

Representative U.S. Patents on Expert Systems, Neural Networks and Natural Language Systems

I. EXPERT SYSTEMS

	PATENT NO.	INVENTOR	ISSUE DATE (PRIORITY DATE)	ASSIGNEE
A.	4,595,982	BURT	6-17-86 (9-2-83)	-----------
B.	4,648,044	HARDY, et al.	3-3-87 (6-6-84)	TEKNOWLEDGE
C.	4,658,370	ERMAN, et al.	4-14-87 (6-7-84)	TEKNOWLEDGE
D.	4,675,829	CLEMENSON	6-23-87 (7-21-84)	INTELLICORP
E.	4,730,259	GALLANT	3-8-88 (3-1-85)	-----------
F.	4,752,889	RAPPAPORT, et al.	6-21-88 (8-16-86)	NEURON DATA
G.	4,803,642	MURANAGA	2-7-88 (7-21-86)	TOSHIBA
H.	4,815,005	OYANAGI, et al.	3-21-89 (11-29-86)	TOSHIBA
I.	4,829,426	BURT	5-9-89 (9-2-83)	COGENSYS
J.	4,866,634	REBOH, et al.	9-12-89 (8-10-87)	SYNTELLIGENCE
K.	4,866,635	KAHN, et al.	9-12-89 (10-19-87)	CARNEGIE GROUP
L.	4,891,766	DERR, et al.	1-2-90 (6-15-87)	IBM
M.	4,916,633	TYCHONIEVICH, et al.	4-10-90 (8-16-85)	WANG
N.	4,928,236	TANAKA, et al.	5-22-90 (7-22-87)	SHARP
O.	4,931,951	MURAI, et al.	6-5-90 (5-8-87)	MITSUBISHI
P.	4,939,680	YOSHIDA	7-3-90 (2-15-88)	HITACHI

II. NEURAL NETWORKS - HARDWARE

A.	4,660,166	HOPFIELD	4-21-87 (1-22-85)	BELL LABS (AT&T)/ CAL TECH
B.	4,719,591	HOPFIELD, et al.	1-12-88 (11-7-85)	AT&T, AT&T BELL LABS/CAL TECH
C.	4,951,239	ANDES, et al.	8-21-90 (10-27-88)	U.S.A. (NAVY)
D.	4,988,891	MASHIKO	1-29-91 (5-9-89)	MITSUBISHI

III. NEURAL NETWORKS - SOFTWARE

A.	4,803,736	GROSSBERG, et al.	2-7-89 (11-27-85)	BOSTON UNIV.
B.	4,884,216	KUPERSTEIN	11-28-89 (11-9-87)	------------
C.	4,897,811	SCOFIELD	1-30-90 (1-19-88)	NESTOR
D.	4,912,653	WOOD	3-27-90 (12-14-88)	GTE LABS
E.	4,941,122	WEIDEMAN	7-10-90 (1-12-89)	RECOGNITION EQUIPMENT
F.	4,954,963	PENZ, et al.	9-4-90 (3-2-89)	TEXAS INSTRUMENTS
G.	4,972,187	WECKER	11-20-90 (6-27-89)	DIGITAL EQUIPMENT
H.	4,979,126	PAO, et al.	12-18-90 (3-30-88)	AI WARE

IV. NATURAL LANGUAGE SYSTEMS

A.	4,670,848	SCHRAMM	6-2-87 (4-10-85)	STD. SYSTEMS
B.	4,829,423	TENNANT, et al.	5-9-89 (1-28-83)	TEXAS INSTRUMENTS
C.	4,916,614	KAJI, et al.	4-10-90 (11-25-86)	HITACHI
D.	4,920,499	SKEIRIK	4-24-90 (9-30-87)	DU PONT
E.	4,974,191	AMIRGHODSI, et al.	11-27-90 (7-31-87)	SYNTELLECT SOFTWARE

INTELLECTUAL PROPERTY LAW ASPECTS OF ARTIFICIAL INTELLIGENCE

by

Thomas K. Dreier
Staff Member
Max-Planck-Institute for Foreign and International
Patent, Copyright, and Competition Law
Munich, Germany

I. Artificial Intelligence and Intellectual Property

Artificial Intelligence--or, in its abbreviated form, AI: the name designates a conglomeration of different programming activities and goals pursued. Today, it includes expert systems, natural language and visual image recognition, artificial speech, controlled motion and learning. In the immediate future, other applications will most likely be added.[1] One may say that the common characteristic of AI applications is to--at least try to--simulate intelligent human activities. The label AI, however, is not without ambiguity. First, no real consensus has so far been reached as to how to define human "intelligence"; second, the question remains as to how an artificial device would have to "behave" in order to be called intelligent.[2] Consequently, it has often been pointed out that the term "AI" is rather misleading and gives rise to false expectations as well as to exaggerated feelings of anxiety.

AI calls for interdisciplinary research. This is true with regard to both the development of each single AI application and to the question of how to deal with the subject itself; and dealing with it includes considering the intellectual property aspects of AI. Assuming that interdisciplinary research has the advantage of bringing together the necessarily limited views of each of the specialists participating, it quite often is impeded by the fact that appropriate solutions cannot, by definition, be provided by an expert in one single field.

This gives rise to two introductory comments. First, as far as expert systems are concerned--the most prominent and until now the most widely marketed of the AI applications--, an excellent analysis based on technical understanding has already been undertaken by Spoor.[3] It would not make much sense to just duplicate his views; and as he has rightly put it himself, "definite answers can only be given by analyzing ... a specific"[4] expert system. Second, working with and representing an institution the task of which is to engage in what would best be termed "fundamental legal research," I must admit that at the beginning of this interdisciplinary dialogue organized by WIPO, and thus in the absence of any further technical information, I do not see myself in a position to adequately analyse other AI applications, or to construct a consistent intellectual property theory in the AI field.

The following discussion paper, therefore, is not intended to be a classical analysis of the intellectual property issues of each single AI application. Rather, it contains some general remarks on the framework within which solutions to the practical and legal problems caused by AI will have to be found, exemplified by--in some cases additional--thoughts on illustrative single issues.

II. Copyright and New Technologies

It should be recalled that the first question is always whether copyright should be open to grant protection to the results of new technologies, or whether copyright should be limited to "traditional" literary and artistic works.

It is claimed here that the answer to the latter question should be negative. First, under the Berne Convention (BC) as well as under national copyright laws, the list of protectable subject matter is non-exhaustive; therefore any "production[.] in the literary, scientific and artistic domain" may be protected (Art. 2(1) BC). Furthermore, historically speaking, copyright has evolved to cover subject matter which at the very beginning was not considered to be within its scope. Thus, the aspects of functionality as well as of industrial manufacturing and application are no longer absolute bars to copyrightabilty in the case of works of applied art; the objection that something resulting from a machine could never be regarded as a "work" within the meaning of copyright has been overcome in the case of photography; and, more recently, in order to bring computer programs within the scope of copyright protection, it has been accepted that the form of a work need not be aesthetic in nature, i.e. neither beautiful nor appealing to the human eye, and, moreover, that it need not even be directed to the human senses as long as it has a defined structure and can be made perceptible.

It should, however, be noted that each of these evolutions has left a trace still visible in copyright law. Art. 2(7) BC leaves it a matter for legislation in the countries of the Union to determine the extent of the application of their copyright laws to works of applied art. In some countries, namely in Scandinavia,[5] photographic works even today receive legal protection only outside the framework of traditional copyright laws. Finally, it hardly needs to be stressed that copyright protection for computer programs has sometimes led to the adoption of specially-tailored rules and that the debate whether or not computer programs are to be regarded as suitable subject matter for copyright protection has still not come to an end.

Moreover, the concept of "human" authorship itself is challenged by some as unduly restricting the scope of copyright. It is argued that copyright protection--as opposed to a neighboring rights or a sui generis protection--should also be available to non-creative organizational and investment-intensive efforts. Furthermore, it is claimed that since it is the investment and therefore the investor who ultimately is in need of and merits protection, this protection should, *ab initio*, be conferred on the natural person or legal entity which provided for the investment and which undertook the economic risk. This, it is said, should be the case notwithstanding the fact that the object of the investment was a human creative activity resulting in a copyrightable work.

The debate on the proper role of copyright in the changing reality of human creativity and technological production is far from over. Yet along which lines should one argue and on what basis are the answers to be found?

III. "Conceptual" Deductive Approach?

It is submitted for discussion that a traditional "conceptual" way of deducting solutions from existing copyright law definitions is a rather questionable if not unsuitable way of proceeding. This does, however, not preclude that in some cases such a way of proceeding might still lead to predictable and hence satisfying results. However, more often than not it seems that new technological subject matter--and this is certainly true for AI applications and their uses--displays neither a structure nor characteristics similar to those of the traditional works which resulted in the formation of some of the central existing copyright concepts.

An illustrative example is the concept of "reproduction." Defined within the context of the print medium, the term could still cope with new phenomena such as reprography and audiovisual copying. We usually know when we are faced with a reproduction, when to prevent it and why. In essence, the political discussion bears on the extent of possible exceptions and on the proper way to police illicit uses. However, what does amount to reproduction when a computer is being used? Of course, it might seem that the storage of a computer program for any not insignificant period of time would fulfill the necessary criteria of both fixation and permanence. However, the purpose of the reproduction right, to enable the right-holder to authorize and thereby receive the benefits of a "normal exploitation of the work" (Art. 9(2) BC), already casts some doubts: does it make a difference whether the storage is on hard disk or on a diskette; whether it is total or partial; whether it is for a longer or a shorter period; whether it is made for duplication, for back-up or for archival purposes; or whether it takes place in the working memory while the program is running? It is not without reason that the term "technical reproduction" has been resorted to more than once in order to describe something which, although satisfying the copyright definition of reproduction, shall nevertheless not be regarded as a reproduction for copyright purposes. Needles to say that the criteria for such a decision lie well outside the concept of reproduction.[6]

The unsuitability of a merely deductive approach to copyright protection becomes even more apparent under German law, which does not treat reproduction and adaptation/translation equally; there is a flat prohibition on the reproduction of computer programs, whereas any adaptation/translation is allowed without the consent of the rightholder, provided the adaptation/translation is not published or exploited.[7] Who could now tell whether--from a conceptual standpoint--the act of decompiling an object code would have to be regarded as an act of reproduction or as one of adaptation/translation?

In this case, any deductive way of reasoning with regard to the elements of a provision, embarked on just to obtain a logically verifiable legal result, would almost invariably and to a great extent be influenced by the desirability of the legal result itself. It need not be emphasized that such a way of defining the elements of copyright provisions would bring with it the risk of interest-biased[8] legal solutions.

IV. Parameters for the Finding of Solutions

If the "conceptual" way of deduction does not bring about the legal results required, what then might be the appropriate parameters for the finding of solutions?

First, the traditional antagonistic values would be legal certainty on the one hand, and individual justice on the other. In addition, adequate and fair solutions might be considered as being opposed to the coherence of the copyright system. "Coherence" would include both the consistency of the set of copyright provisions and the adequate treatment of both the creators of works traditionally protected by copyright, and of those who create new protectable subject matter.

Second, monopolistic protection in favour of an individual creator may be regarded as being contrary to free access to protected subject matter by the public at large. However, both aims may also be seen as being linked together since one of the main reasons for granting monopolistic protection is to enhance creation and therefore create works for the benefit of society.

Third, another decisive factor for the finding of a solution has been whether the advantages of the established system of protection under the international Conventions outweigh the disadvantages of bringing new subject matter into the traditional copyright framework rather than drafting sui generis rules. As the case of computer programs on the one hand, and of integrated circuits on the other hand has shown, the answer may not always be identical.[9]

First and foremost, however,--and this is put up for discussion--rather than looking at the law and deducing legal solutions it might contain, the problem-finding mechanism, somewhat simplified, would have to work "backwards": it would have to start with the definition of adequate and acceptable solutions, and see to what extent they fit into the existing system of intellectual property protection. Insofar as they do not correspond, special legislation would have to be drafted, or the solutions to be adopted would have to be modified so that they do come within the scope of copyright.

As it may seem, more often than not this is the process already applied in search for legal solutions to the protection of new subject matter. The elaboration of the draft EC Directive on the legal protection of computer programs may be perceived as an almost perfect example for this way of solving the problem. Without going into detail, it may be said that in the end[10] priority has largely been given to the coherence of the copyright system. Moreover, it became apparent that the different paramaters have been used in order to serve the individual economic interests of the parties concerned, rather than themselves forming a coherent theory.

With this in mind, some remarks on single problems within the field of AI applications may now be looked at. In a way, the following illustrative issues have as little in common as the different AI applications.

V. Some Examples

Intellectual property issues arising from AI may be spotted in each of the three fields of copyrightability, ownership attribution and infringement.

A. Questions of copyrightability

The following remarks shall be directed to three problem areas: collections of data, facts and rules; algorithms; and computer-output.

1. Collection of Data, Facts and Rules

Knowledge bases of expert systems consist of a body of facts and rules. Similarly, data banks work on the basis of information, data or any other subject matter which can be electronically stored.[11] Of course, any item contained in a knowledge base or stored in a data base may itself be copyrightable and its storage therefore be subject to the respective rightholder's consent. In principle, this does not seem to pose a major problem. What will be considered here, therefore, is the question whether, and if so under what conditions the collection of such items--especially if they constitute what might be called "works of fact"[12]--does or should enjoy intellectual property protection.

As an arrangement of protected and/or unprotected subject matter, a collection might be protected by copyright, provided the criterion of originality is fulfilled. However, this might be problematic under all those national copyright laws requiring more than the mere fact that the collection must be from its author and must not be copied. Originality presupposes possibilities of selection. Only creative choices made justify the protection conferred, not effort or investment alone.

It would, however, seem that the creative choices to be made when assembling data for a data base are rather limited.[13] In essence, the selection would consist in choosing in which field to collect the data and the data of which fields can be interrelated by the prospective user of the data base. Take, as an admittedly simple but nevertheless instructive example, the case of a computerized address directory listing names, street addresses and, let us say, professions. Since the data as such are free and non-copyrightable and therefore cannot contribute to the originality of the collection, the only creative element in this case would be in the selection of exactly these data fields (and not others like, e.g., the size of the families or the number of cars etc.).[14] Furthermore, there would be no selection of the data within this field if the goal to be achieved was completeness. Finally, there seems to be little if no arrangement of the data within the data base memory, since the data must be equally accessible and since any order or arrangement is only achieved by the software of the data base according to the specific search instructions given by the end user. Of course, given the relatively limited importance of protection of collections until now, only little case law seems to exist. While some national decisions do affirm that protection might exist as long as the "selection and arrangement" show sufficient creative choices,[15] to merely choose a certain field and then collect all the data therein would run a relatively high risk of not being sufficiently original to attract copyright protection.[16]

This analysis might produce different results with regard to the information stored in the knowledge base of an expert system. It is true that as far as the knowledge base contains general rules, these rules themselves could not be copyrightable, nor could any data, facts or information as such. However, the activity of a knowledge engineer seems to be much more complex than just making minimal choices with regard to what information to store and what information not to store in the knowledge base.[17] First, when interviewing an expert, the knowledge engineer already makes a selection by asking certain questions and omitting others;[18] second, he selects and transforms the answers; and third, to some extent he may also have to arrange the material for storage in different parts of the knowledge base. It should be taken into account that some of the selection and arrangement closely follows the structure of the respective field of knowledge, and that one of the parameters for selecting might be completeness of the information and expertise stored, so that any such selection would seem to be dictated by function and could, therefore, not contribute to the originality of the knowledge engineer's activity. However, in general there might nevertheless be a good argument indicating that there would be enough room for the knowledge engineer to make creative choices.[19]

When drawing the line in an individual case, the dual function of the originality criterion should always be kept in mind. Not only does originality retrospectively indicate which products may enjoy copyright protection and which may not, i.e. by comparing the object in question with prior works in the same field. Likewise, it prospectively determines the extent of the material being protected, which may not be taken by future creators without authorization.

To return to the example of the address directory listing names, street addresses and professions: if copyright protection were granted on the basis of and for the choice of these three fields of data only, this would invariably have the consequence that nobody else could compile a directory based on the same choice of criteria;[20] and this although the data as such are undisputedly free for everyone. The only exception might be an independent creation, which is permissible under copyright law. However, the independency would have to be ascertained not in the independence of the realization of the second work but in its creative elements, i.e. in the selection of the respective data fields. Since subconscious copying will also constitute an infringement, cases of independent creation will be extremely rare.

In sum, any protection granted on the basis of only a minimum of originality and therefore to little more than a mere idea would in fact extend the monopoly unjustifiably to non-copyrightable material and thus exclude third parties from entering the market using the same unprotected information in the same or a similar field, and would therefore severly hinder competition.

Of course, any neighboring rights or unfair competition approach could easily avoid this pitfall,[21] since it might not protect the collection as such against any act of reproduction but only against any unfair appropriation such as outright copying or any partial taking with changes insignificant from the point of view of the effort and investment made, i.e. the real object of protection.

2. Algorithms

Any AI application is based on, or contains a computer program which itself is based on one or, most likely, on several algorithms. While it may be stated that copyright for data-processing programs is now widely accepted--since as a result of the technical development, the intellectual property protection of computer programs has first become a matter of interest--so that little needs to be added here, the (non)protectability of algorithms seems to have received much less attention so far. This gives rise to some remarks and open questions.

An algorithm may be defined as a general procedure to solve, in a finite number of fixed steps, a given problem out of a defined class of problems. As a general mathematical rule, the algorithm is denied protection, since it is said that as such it must remain free to be used by everyone and cannot be monopolized.

Therefore, whenever the copyrightability of anything containing or looking like an algorithm is being discussed, the legal analysis does not so much focus on what constitutes an algorithm, but on whether or not it should be free from protection and available to everyone. In essence, this is another example of the "solution-oriented" (as opposed to the "conceptual"[22]) approach to finding legal solutions for new subject matter.

However, the rule that any algorithm is denied protection is not as clear cut as it may seem. In some countries at least, algorithms may be patented, provided the invention as claimed does not encompass any and every means for performing a certain function.[23] Would it not seem possible to similarly grant copyright protection at least to relatively complex algorithms which, in fact may be regarded as combinations of several algorithms? It takes quite considerable personal and financial resources to develop certain algorithms. Are they not in need of protection?

On the one hand, it may be argued that copyright protection, although conferred without formalities, and not being an absolute monopoly right[24] would, in the case of algorithms, in fact extend far beyond the scope of protection available under patent law, since it would cover any possible use of the algorithm. Moreover, it might be claimed that there is no real need for the long term of copyright protection, since most complex algorithms will most likely have been developed to be implemented not by <u>any</u> but only by <u>specific</u> means, so that patent protection would be available.

On the other hand, does the combination of algorithms in some instances not already form the structure of a computer program? Structure, sequence and organization, however, have been held to contribute to the originality of protectable subject matter. The question would then be, when are we still faced with one algorithm, and when with a combination of several algorithms? Here, the answers given by programmers seem to indicate that a clear dividing line does not exist and cannot always be drawn. This, however, might lead to the conclusion that perhaps the algorithm should not be excluded per se from copyright protection and that the answer might rather depend on how far the monopolization in an individual case would actually extend.

Nothing of this, however, should be taken as a definite answer; rather it should be seen as a proposal to reconsider the matter carefully.

3. Computer Output

Finally, questions of copyrightability arise with regard to the output of AI applications. Whenever a computer is involved in shaping the actual form of the output, the very concept of <u>human</u> authorship might seem questionable.

Of course, the output of AI applications may occur in different forms. It may have been reduced to a tangible form such as a printout or an object shaped by computer-operated tools. It may also be intangible, e.g., information appearing on the screen, or a certain order of objects to be positioned or conditions and states to be controlled. Admittedly, copyrightability of output in the last category seems rather doubtful, since probably it would not take the form of a "work" within the meaning of copyright laws.[25] However, the question whether or not the output is in a tangible or in an intangible form--e.g., whether, and under what conditions something appearing on the screen may enjoy copyright protection--would not as such be relevant to the copyrightability issue. Rather, copyrightability would depend on the fixation requirement of national copyright legislation and therefore cannot be dealt with in detail.

The sole question to be treated here shall be under what conditions output which has--or which may be reduced to--the form of works traditionally protected by copyright will be regarded as copyrightable subject matter, despite the fact that a computer has been involved in the shaping of its actual form.

Two basic assumptions seem to be in conflict. The first is that such output may no longer "bear the mark of the personality" of the person responsible for the corresponding input, since the actual form has been shaped by a machine and not by a human being. The other assumption is that it would be economically unwise and even counterproductive to exclude certain human-initiated machine-output just because a computer has been used as the best or even the only machine able to achieve a specific result.

More inclined towards the second of these two assumptions, the following concept is put up for discussion here:

First, the mere fact that a computer has been used in shaping the form of an object which from its appearence would qualify as a "work" within the meaning of copyright, should not per se exclude this output from copyright protection.

Second, a distinction would have to be made between, on the one hand, output during the creation of which the computer has merely been used as a tool in order to help realizing the form of an already preconceived idea, and, on the other hand, output during the creation of which the computer has been used as an instrument in order to conceive--at least parts of--the idea itself.

The first category--including, e.g., text processing programs, paint-boxes and various computer-aided applications--does not pose a particular problem, since indeed there seems to be no difference in comparison with traditional tools such as pen and ink, a chisel or a brush. The second category, however, where the structure of a work is determined, e.g., by a random device, or where the computer program has itself performed the main choices essential for the structure or the form of the output,[26] needs closer examination. Here, the question arises as to what extent human input activity might be sufficient to meaningfully attribute the output determined by this activity to the person(s) responsible for the input, and therefore treat the output as copyrightable subject matter.

On the assumption that the search for human authorship should not be given up all too easily,[27] it should therefore, third, be retained that human input activity might be sufficient to find creative activity for the purposes of copyrightability.

And fourth, as a further distinction it is proposed that sufficient originality should be ascertained in the input activity—of either the programmer and/or the user[28]—as long as the person(s) responsible for the input at least decide(s) on the framework or the basic parameters of the actual output and if the input activity leaves sufficient room for creative choices. Since it is inherent to the possibilities of a computer to produce whole series of output, the creative act would necessarily also include selecting activities, insofar one might speak of computer-aided works. Only where there is little or no human input from the user's side and where, at the same time, the activity of the programmer is all too remotely linked to the actual form of the output, should the output not be regarded as copyrightable subject matter. This seems logical as long as one assumes that authorship presupposes some kind of human activity, and since therefore the computer as such cannot be regarded as "author" within the meaning of copyright. Such output might then be termed computer-generated.[29] It may be noted that the U.K. Copyright, Designs and Patents Act 1988, the first legislative attempt at dealing with the issue, seems to operate a similar distinction by defining a "work generated by computer" as one generated "in circumstances such that there is no human author."[30]

In sum, the result would be that computer-aided works, as defined here, would constitute copyrightable subject matter, whereas computer-generated products would not. In order to protect the latter as well, any legislature would, of course, be free to enact legislation to fill this gap. This might be done by a specifically tailored neighbouring right, or under an unfair competition law theory which could then protect the person having undertaken the respective effort and/or investment. It should be added that the U.K., in line with its property-oriented understanding of copyright,[31] did fill the gap despite the absence of human authorship by granting a copyright to "the person by whom the arrangements necessary for the creation of the work are undertaken."[32]

It is true that in practice, the line between computer-aided works and computer-generated output may not always be easily drawn. To give just one example, the use of a random device would not per se prevent creative input activity being ascertained. The task is then to determine, in an individual case, whether or not the possibilities of creative and not function--dictated input choices, and the way in which these choices have been made, would reveal sufficient creativity. However, this should not be any more difficult and should not lead to less certain legal results than judging the originality of works created by traditional means.

B. Attribution of ownership

In essence, four remarks seem to be called for.

First, as has already been pointed out, whenever national legislature decides to fill the gap and protect computer-generated products, he would also be free in his choice of to whom to attribute the ownership of this newly created right. Most likely, he might chose the person who undertook the production effort or made the necessary investment.[33]

Second, as far as computer-aided works, as defined above, are concerned, the traditional rules of authorship attribution would seem to apply without great difficulty. If several persons have made creative input, it all depends on their respective contribution.

At one end of the scale, when a computer is being used as a tool,[34] all the creative choices are being made by the user. Thus, nobody would seriously claim that the programmer of a text processing program has any rights in the texts written by the user with the aid of the text processing program. The same would be true of any graphic work designed with the aid of a paint box.

At the other end of the scale, we would find cases where the user does little more than press some keys. The menu-screen or user-interface, to cite an example, he thus activates has as such already been determined by the program and therefore would be the programmer's[35] exclusive property.

As far as the continuum between these two ends is concerned, one does not necessarily have to assume authorship of both the programmer and the user whenever there is some choice left to the user. Rather than being a question of mechanical logic, any decision will be based on an assessment of the creative input contributions of the persons responsible for the actual form of the computer output. Thus, the courts have held that although the player of a video game is left with some choice of how to react, this does not make him a co-owner of the game, since the player is said to stay within the limits of what is predetermined in the game program. Conceptually speaking, this would also be the case when using a paint-box program, since the user cannot produce brush strokes outside the parameters predetermined by the program. However, there is no doubt that the user would be the sole author of any paint-box-aided painting.

Third, whether co-authorship will be ascertained in an AI application, or whether the application will have to be looked upon as a derivative work or as any other form of joint authorship, largely depends on how these different forms of ownership are defined and what intensity of a common proceeding is required by the respective national legislation.[36]

The same would be true with regard to the question of whether the copyright of AI applications may originally vest in a legal person and, likewise, whether it may vest in a person different from the natural person actually having made the creative contribution. It is well known that the prevailing view--according to which the term "author" as used in the Berne Convention would only allow for an assignment or transfer of rights, with the possible exception of cinematographic works[37]--has come under attack in the international debate.

Fourth, it has been pointed out that different legal regimes might be applied to the individual components of an AI application such as expert systems. In particular, this seems to be the case in France, where the legislature in 1985 decided to draft sui-generis rules, albeit within the larger framework of droit d'auteur, for the protection of computer software; whereas the intellectual property aspects of the knowledge base would be governed by the general Act on Literary and Artistic Property of 1957 as modified.[38]

However, in the end this problem does not seem to be much more complicated than where several persons have contributed to a traditional work which is being marketed as one. Its components may have been created independently and only been joined together later on as in the case of, e.g., empty software shells. On the other hand, as is probably true in the case of most existing AI applications, they may have been developed with the outright intention of being put together. In general, copyright law does provide at least a framework of solutions which would be further defined by case law. Where sui generis rules do not provide such a framework, they may be presumed to have been all too hastily drafted, and future legislation should try to avoid such pitfalls.

Furthermore, any discrepancies in legal treatment, as well as any problems resulting from joint or co-authorship might be solved--at least to a large extent--by contractual arrangements under existing copyright legislation. This seems all the more true, since anyone who wants to market AI applications should be in a position to know what rights to which components he or she would have to secure from which persons.

C. Infringement problems

As far as infringement is concerned, the legal analysis would seem to follow known routes to an even greater extent. Cases of infringement can be found at all levels: when storing protected subject matter in a data bank or a knowledge base; when adapting, translating or otherwise changing protected material within an AI application, generally provided that such adaptation will be made available to the public; when receiving, retrieving or otherwise obtaining the results of an AI apllication; and, finally, when taking part or all of a protected AI application itself.

Some issues, however, might call for closer examination.

One such issue would be the moral rights aspect. Any work protected, once it has been digitally stored, can easily be duplicated and manipulated with an ever increasing range of software. It may be altered, cut, combined, colored, rotated, enlarged or otherwise contorted. If such acts do constitute an infringement of moral rights and/or of the adaptation right, it may seem difficult to determine how such acts, if unauthorized, could be monitored and prevented by the individual right owner. So far, photographers seem to become the first victims of the new technical possibilities.

A further problem would be partial taking. Since the copyright standard is whether the allegedly infringing work is substantially similar to, or, as the national formula may read, not substantially different from the protected original, something less than the protected work or any unprotected portions thereof, may be freely taken.[39] Of course, this is nothing new and it remains to be seen whether and to what extent this phenomenon would arise and what its economic consequences would be in the field of AI applications.

As far as the user side of AI applications is concerned, it might at first sight not seem clear which acts would be subject to the authorization right of the right holder.[40] For instance, would on-line retrieval need authorization, and would it make a difference whether or not the user makes any printout and if so how many? If on-line retrieval would possibly not amount to a reproduction, might it be considered a communication to the public? A further difficulty arises in the fact that the user is not always interested in a protected work but more often than not he will be seeking the information it contains. In this respect, a user may well be satisfied to receive the answer that no information is available on the question he has asked, and this transfer of information does not involve the communication of protected subject matter.[41] Furthermore, each single appropriation of protected material by the users may be de minimis and difficult to police, but taken in total may "seriously erode economic and moral rights in original works."[42] However, it should not be impossible to find appropriate and adequate solutions, comparable to those already developed with regard to other mass uses of protected material such as reprography and home taping, which are difficult to police but would nevertheless unduly infringe upon the authors' rights.

It should, of course, not be forgotten, that as far as the appropriation of protected material for the purpose of an AI application is concerned, a good deal of issues might successfully be taken care of by way of contractual agreements. This would also be true with regard to the relationship between the authors of AI application components, AI application producers, sellers and/or those offering AI application services and the ultimate users.

What remains is the protection of the AI application as such against any taking in whole or in part by a competitor. In cases where there is no copyright protection available to the AI application, and where the appropriation cannot be prevented by other means, such as unfair competition law remedies, but where the taking would nevertheless be unfair and economically undesirable, legislation is indeed called for.

VI. Conclusion

One of the major problems in bringing investment-intensive AI applications into the framework of existing copyright laws seems to be that the latter protect the creative work result rather than the effort or investment which has gone into a product. At best, effort and investment might serve as indicia for a sufficient degree of creative activity.

However, once the issue of the copyrightability of technologically new subject matter has been solved, the application of traditional copyright rules to a given issue within the field of AI applications will cause much less problems than one initially might have expected. When such problems arise, they will quite likely be not much more complicated than any other issue under traditional copyright law, provided the technical facts have been made sufficiently clear to the non-expert lawyer.

Possibly, the question might arise whether or not a certain act of appropriation of protected material would be covered by the law or not. It seems inevitable that the law, in order to develop and keep pace with technological progress, will have to work on the basis of analogy. This it has always done. The present economically-oriented way of legal thinking in the field of intellectual property protection seems to suggest a "solution-oriented" rather than a more traditional "conceptual" approach. It is for the courts to decide when the limits of judging by analogy will be reached, yet it is a political issue for the legislature to decide when and how to draft sui-generis rules, as well as how to police certain uses of protectable subject matter within the field of AI applications.

NOTES

1 For possible future applications see, e.g., Brand, The Media Lab--Inventing the Future at M.I.T., New York 1987.

2 According to the Turing test, a computer would perform intelligently, whenever the person communicating with it could not distinguish its answers from the answers given by a human being. However, from a philosophical point of view it might be asked whether, in order to be regarded as "intelligent," the working of a computer would not have to be meaningful in itself, or whether beyond this a computer would have to be required to be conscious of its own working; see, e.g., Hofstadter, Gödel, Escher, Bach: An Eternal Golden Braid, New York 1979.

3 Expert Systems and the Law--An Outline, 1990 Copyright 330 et seq., with further bibliographical references. Literature on computer-aided and/or computer-generated output as such is now more readily available. See, e.g. Samuelson, Allocating Ownership Rights in Computer-Generated Works, Univ. of Pittsburgh L. Rev., Vol. 47, No. 4, pp.1185 et seq. (1986); Goldstein, Summary Report, as well as national reports in: ALAI Canada (éd.), L'informatique et le droit d'auteur--Actes du 57e Congrès de l'ALAI, Montréal 1990, pp. 433 et seq.; Sookman, Computer-Assisted Creation of Works Protected by Copyright, 5 Intellectual Property Journal 165 (1990), pp. 179 et seq.

4 Spoor, op. cit., at 333.

5 See, e.g., Karnell, Photography--A Stepchild of International Conventions and National Laws on Copyright, 1988 Copyright 132 et seq.

6 The draft EC directive on the legal protection of computer programs has met with the same problem. If in its present version (COM (90) 509 final, O.J. C 330 of 20 December 1990, pp. 22 et seq.) it considers "loading, viewing, running, transmission or storage" of the programs as restricted acts "[i]nsofar as they necessitate a reproduction of the program in part or in whole," the term "reproduction" rather than being defined has been referred to without any further clarification. Needless to say that such a rule merely restates what has already been known, still leaving the problem unresolved.

7 Secs. 53(4) and 23 of the German Copyright Act.

8 In the sense of interests influencing the solutions which, since they operate on the level of the definition of the elements of a rule, have not been laid open, as opposed to an interest-influenced way of finding the solution which clearly states on what particular interest the reasoning finally retained is being based upon.

9 For a comparative study see Dreier, National Treatment, Reciprocity and Retorsion--The Case of Computer Programs and Integrated Circuits, in: Beier/Schricker (ed.), GATT or WIPO? New Ways in the International Protection of Intellectual Property, Weinheim 1989, pp. 63 et seq.

10 I.e., at the stage of the Common Position adopted by the Council pursuant to Art. 100a, 147 of the EEC Treaty on December 13, 1990.

11 This may be any object which already comes in, or which may be reduced to digitalized form. Consequently, in order to enhance the possibilities of future information exchange, digitalization of fields as divergent as broadcasting and film, printing and publishing, as well as computers and telecommunications has become a major goal to be achieved within the not too distant future; see, e.g., Brand, op. cit., pp. 9 and 201 et seq.

12 Cf. the proceedings of the "Copyright in Information" conference 1989, edited by Dommering/Hugenholtz under the title "Protecting Works of Fact," Amsterdam 1991.

13 This has rather pointedly been shown by Metaxas, Protection of Databases: Quietly Steering in the Wrong Direction, (1990) EIPR 227 et seq.

14 Similarly, in Feist Publications Inc. v. Rural Telephone Service Co. Inc., the owner asserting a copyright in his telephone directory argued before the Supreme Court that the following choices had been made: "(1) to put out one phone book as opposed to 26 separate books; (2) to include business listings; (3) to include the full address in some listings and just the town for others; and (4) in what format to publish the names"; see 41 PTCJ 247, at 248.

15 A typical example would be the selection and arrangement of texts of world literature in an anthology; see for German law Federal Supreme Court of September 25, 1953, GRUR 1954, 129 et seq. For further references to German law see Schricker/Loewenheim, Urheberrecht, Munich 1987, /2 note 61, and /4 notes 3, 7 and 9. However, contrary to early cases considering phone directories copyrightable, the mere alphabetical arrangement and simple classification usually used in similar products has been held not to be sufficient; see RG JW 1925, 2777, and BGH, GRUR 1961, 631 at 633; see also more recently BGH, GRUR 1987, 704 et seq., concerning a dictionary of trademarks.

16 Admittedly, this leads to the question of how to define the area within which to look for choices possible in order to judge the creative character of the choices actually made. In general, this seems to be the area described by similar or corresponding preexisting works. If so far no comparable preexisting work has been created, one might either consider the discovery of a new area as contributing to the originality of the newly created work, or define a larger area, or compare the work in question with similar works conceivable or expected within the same new area.

17 For the activity of a knowledge engineer see, e.g., Karbach/Linster, Wissensakquisition für Expertensysteme, München 1990, pp. 21 et seq.

18 Depending on the technique of the interview, the questions asked most likely will in turn be influenced by the answers given by the expert interviewed. This, however, would influence the question of ownership (see below), rather than the question of copyrightability.

19 See, as an earlier, but relatively close exemple, the German decision of the LG Düsseldorf of February 2, 1968, in: Schulze, Rechtsprechung zum Urheberrecht, Looseleaf, No. LGZ 104, which held that "the determination and judgment of those accounting cases which an accountant has to know in order to meet with the requirements of normal accounting" show sufficient originality for copyright protection; ibid, at 5.

20 Of course, given the limited originality, the corresponding scope of protection would have to be rather strictly limited so that in all likelihood the replacement of one of these three data fields by another would not amount to infringement. Similarly, the taking of the data of one field only would be nothing more than the taking of unprotected material.

21 Consequently, the problem of copyrightability is much more acute in countries where there is no--or not a sufficient--unfair competition law. Thus, in the U.S., if trade secret protection should not apply, the only alternative seems to be an overbroad protection by copyright or no protection at all; cf., e.g., Judge O'Connor's remark at the hearing of the Feist case: "[I]f .. respondent's white pages contain no copyrightable material .. therefore anyone could just copy the pages and sell them?"; 41 PTCJ 247.-- Indeed, the Supreme Court, in its decision, has held the collection of names in the white pages to be non-copyrightable material; see Decision No. 89-1909, March 27, 1991, 41 PTCJ 453.

22 For discussion see infra.

23 For U.S. law, see especially In re Grams, 12 USPQ 2d 1824, and In re Iwahashi, 888 F.2d 1370, the latter invention claiming an autocorrelation circuit for use in pattern recognition.

24 However, it must remain doubtful how much freedom the possibility of an independent creation existing under copyright, but not under a truly exclusive right such as a patent, would actually leave in the field of finding and formulating algorithms.

25 However, this is not as clear as it may seem at first sight. Whereas on the one hand the controlled state of, e.g., physical objects such as the temperature of a liquid would indeed hardly be a "work," the positioning of objects in space on the other hand might qualify, depending on whether the plan of the positions would be regarded as copyrightable subject matter or not. Thus, an AI application which would put blue round devices into the left, and red square devices into the right box, would not produce output which might qualify for copyright protection, whereas the output an AI application which forms a certain pattern out of these two components perhaps might.

26 In reality, however, this line is sometimes rather unclearly drawn. Especially, a proper understanding of the tool qualities of a computer would mean that the computer is used whenever the effect desired could better or even only be achieved with its help. Thus, the mere use of a computer as a tool may well influence not only the form but to some extent also the essence of the work. But this is nothing new, since already the choice by the creator of traditional tools such as a particular pen or brush is motivated by the very characteristics of their respective trace.

27 It is admitted that in general, this assumption might be debatable; however, this cannot be deepened here, since such a debate would seem to be linked to the much wider discussion whether, in order to best serve the purpose of furthering creativity, protection should be vested with the individual author, or, rather whether it should be conferred on the producer as the person having to directly deal with the actual users of protected subject matter.

28 Whose input activity would show creative elements would be a question of ownership attribution; see text following.

29 It would seem advisable to avoid the term computer-generated work, since it might result in the misleading impression that a product, the essential structure and form of which have been determined by the working of a computer, would be a "work" within the meaning of copyright, and the computer its "author." Similarly, to speak of a computer-generated work also in cases where a sufficient degree of creative input activity may be ascertained, tends to obscure the fact that it is the person responsible for the input activity and not the computer who would have to be regarded the author of the resulting work.

30 Sec. 178 of the Copyright, Designs and Patents Act 1988 (CDPA).--However, the CDPA, despite the absence of human authorship, nevertheless speaks of a "work." Also, the CDPA does not warrant criteria to determine under what conditions "there is no human author." Since the provision is modelled after those attributing authorship to the producers of films and of sound recordings--where, however, there is human authorship involved, albeit none of the producers--it might be assumed that absence of human authorship may be found all too easily.

31 See Sec. 1(1) of the CDPA 1988: "Copyright is a property right which subsists ... in the following descriptions of work."

32 Sec. 9(3) CDPA. For ownership attribution, see also text following.

33 For Sec. 9(3) of the U.K. CDPA see text above. It should be mentioned that the EC Commission, in its draft directive on the legal protection of computer programs, and following an amendment proposal by the European Parliament, no longer maintains its former proposal to entitle "the natural or legal person who causes the generation" of programs generated by the use of a computer program, to exercise all rights; see amended proposal, O.J. No. C 320 of December 12, 1990, p. 26.

34 For the distinction between "tool" and "instrument," see text above.

35 Several programmers may, of course, have concurred in the creation of the program.

36 Another question would be how to define co-authorship in the light of ever increasingly facilitated access to digitalized protected material, and how to police the growing number of adaptations; see also text following.

37 Art. 14bis(2)(a) Berne Convention.

38 For details see Chatain, Régime juridique des systèmes experts--Loi de 1957 ou de 1985?, Expertises No. 125 (1990), pp. 54 et seq., who argues in favor of treating expert systems as unitary pieces of software and therefore to only apply the 1985 Act.

39 An already classic example would be the technique of sound sampling. Although it is doubtful that copyright in the work recorded could be used in order to prevent the taking of single sounds or of small portions, since as such they would not show sufficient originality, the protection granted to the recording may be more successful, since it extends to any single part of the recording.

40 Of course, not any act performed in, by or in the course of the working of any AI application with regard to protected subject matter can be discussed here in detail. It is, however, suggested, that a traditional analysis would contribute considerably towards solving these issues, especially if it is kept in mind that exclusive rights, apart from their moral rights aspect, are mainly granted in order to secure for the right holder adequate control of the exploitation of his works.

41 For this example see Metaxas, op. cit., p. 232-3.

42 Goldstein, in: ALAI (ed.), op. cit., p. 450.

INTELLECTUAL PROPERTY ASPECTS OF ARTIFICIAL INTELLIGENCE: SOME ISSUES

by

Ashok Bhojwani
Managing Director
TSG Consultants Limited
New Delhi, India

1. Introduction

1.1. Traditional laws of intellectual property have difficulty coping even with the software paradigm. Additional issues and subtleties arise out of the use and capability of AI systems, namely:

the ability of AI systems to create, adapt, and use for commercial benefit, works covered by intellectual property laws, e.g., inventions, literature, music, art, etc.;

the economic value of knowledge and the ability of AI systems to create new knowledge;

the use of AI systems as tools or assistants for the creation of intellectual properties;

the inclusion of AI technologies within intellectual property, e.g., the inclusion of AI within an information repository or a multi media training course.

1.2. These have an impact on various rights such as economic rights, moral rights, neighboring rights and duration of protection. This also raises issues of machine rights and their relationship to human rights and the need for clarity on the question: should we start protecting intellectual creations like algorithms, ideas, concepts, knowledge and scientific discoveries, which are not traditionally protected under the laws of intellectual property?

1.3. As a result of these capabilities of AI systems, there is a need to reexamine the philosophical foundations of our legal systems and, in particular, intellectual property laws. At the philosophical level there is a need to accept that intelligent machines, which are non human, can be creative and can act independently.

1.4. It may not be possible to redesign the copyright laws to address the consequential issues, without compromising the basic protection for traditional literary and artistic works. Works created by AI systems should either be placed in the public domain or new *sui generis* laws ought to be designed for software and AI systems. These laws should take an integrated view of computer software, hardware, embedded systems, AI technologies and other new developments which are taking place at a fast and furious pace.

2. AI Systems: Scope, Domain and Limitations

2.1. What is AI?

It is difficult to define the term artificial intelligence with any degree of precision. There are philosophical difficulties in defining the terms artificial and intelligence. Nevertheless, it is useful to look at various definitions to try and understand some of the boundaries of the concept of AI.

Some of the textbook definitions [B3, R1, S3] are:

The goal of AI is the construction of hard problems.

AI systems do things at which, at the moment, people are better.

The specific task of an AI system is to interpret its inputs in such a way that its performance gradually improves.

The Association of Computing Machinery (ACM), Special Interest Group on Artificial Intelligence (SIGART) has an interesting way of describing AI:

> AI seeks to understand and apply the principles and mechanisms underlying intelligent behavior. AI hypothesizes that the basic problems of reasoning, representing, perceiving and acting remain constant across hardware that varies between humans and machines, and that viewing these problems as computational tasks offers valuable insights. AI attempts to demonstrate its theories and techniques by constructing machines that can see, hear, speak, read, write, feel, move, make decisions and plans, learn from experience, and exhibit common sense and expertise in solving problems.

Essentially, an AI system is able to manipulate objects through the use of processes. Objects may consist of numbers, text, files, books, paragraphs, musical works, computer programs, etc., and collections of objects. Objects may be organized in the form of structures which are related in some physical way. For example, a paragraph immediately follows another. The processes operate on objects to, inter alia, create other objects (or instances of the same object), modify or reproduce objects, organize objects, and destroy objects.

Thus, an AI system could create paragraphs and organize them in the form of a book. It could then publish the book.

The basic hypothesis[N1] is that such an approach of object manipulation represents the way that human beings behave, and that such a system is capable of general intelligent action.

2.2. Scope of AI Systems

AI has been primarily perceived as a tool for research into problem solving and qualities like intelligence and learning ability. AI holds the promise for better human machine interaction with the machine adapting more to natural human modes of communication--speech, vision, multiple senses, touch, etc.

Research problems that currently fall within the scope of AI[R1] include game playing, theorem proving, general problem solving, perception (vision, speech, etc.), natural language understanding, expert problem solving, symbolic mathematics, medical diagnosis, chemical analysis and engineering design. In addition, some work has been done on the general concept of using an AI system as an assistant for various creative and routine tasks.

Intellectual property creations of AI systems appear to be excluded from this list of research problems. For example, the production of intellectual property like original works in the form of literature, music, art, computer programs, video films, intelligent adaptations, electronic and hard publishing, speech synthesis, transmission and reception over satellite and other networks, generation of non-mathematical ideas, algorithms and concepts.

2.3. AI Methodologies

In an AI system knowledge is represented using approaches like Predicate Calculus, Production Rules, Frames and Scripts and Semantic Networks. These approaches provide varying levels of efficiency and naturalness for different types of problems. These representations are extensions of traditional data structure like lists, tables, hierarchies, sets, rings, networks and matrices used in non AI software.

A computer, coupled to sensors and driven by software, can control movement based on the inputs that are sensed by the sensors. In traditional non AI software, the knowledge about the functions to be performed is hard coded into the computer program. The data on which the program operates is stored on magnetic media or some other device and is referred to as a data base. In an AI system the knowledge is made explicit in the form of production rules or some other representation. The function of the software now becomes one of interpreting the knowledge. The knowledge base, which embodies the knowledge, is separate from the program, is stored on a storage device and may be changed independently of the program. Thus, as rules accumulate, the intelligence of the system improves.

A non AI type of computer program cannot handle situations that it is not programmed to deal with. On the other hand, an AI system can ask that a new rule be given to it so that it may handle a new context. An AI system, typically, can also explain its reasoning so that it may be verified and corrected by a human expert or by its own experience and knowledge.

AI systems are capable of reasoning, inference and deduction. An AI system can communicate in a natural language, for the time being consistent with the semantics of the knowledge domain and at speeds limited by current technology. It can treat the problem domain deeply, flexibly and broadly from several perspectives by subjective and factual reasoning.

2.4. Are There Any Limits to What AI Systems Can Achieve?

The limits achievable by AI systems are not known. HAL 9000, the computer that communicated with humans, in the film 2001--A Space Odyssey, is considered by some researchers[11] as a goal to be achieved, although they do not necessarily like to talk about it, perhaps because of the potential social impact and the fear that robots and AI systems inspire among humans.

AI systems have the potential to change the way that people view their traditional physical and mental activities. This could lead to new states of occupation, leisure, creativity, introspection, gourmandism, drugs, hallucinations, spirituality, etc. Some of these states of being have probably never before been achieved on a mass scale. Some questions that arise:

- How well can we understand the working of our own mind and intellect? How close can we get to the final understanding before we start interfering with the basic nature of what we are trying to understand? Is there a principle similar to Heisenberg's Uncertainty Principle which will come in the way of humans understanding human intelligence completely?

- If there is a limit to humans understanding the nature of human intelligence, will we entrust the task of uncovering the ultimate secrets to a machine?

- What differences will there be between AI and human intelligence? Will there be a difference in the levels of creativity, innovation, recognition, classification, capacity and ability to learn about new objects and situations?

- Will we ever reach a stage of development where machines, or some biotechnology innovations, become more capable than humans? AI systems already seem to be superior in specialized areas needing large amounts of information and drawing rapid inferences from the information.

There are many such questions. All the answers are not there, but they need to evolve. There is a need to philosophize over such issues.

3. General Philosophy

Technology is the means to an end, not an end in itself. It is a tool to achieve the larger goals. What are these goals and how do we arrive at a consensus on them?

Philosophers and religions tell us that our basic goal is to be happy and in total harmony with nature. True happiness comes from within. It is a state of mind and not of materialism. Will AI systems give the potential to take us away from external concerns and allow us the time and the freedom to explore the inner world? Or will they remove the hardships which seem to give us mental strength? Will AI allow us to unleash the innate human potential, and our aspiration for higher values, by removing the constraints which inhibit its realization? Or will there be an increase mainly in hedonism?

There is a need to integrate the philosophies of the East and West and once again seek answers to fundamental questions, with the added assumption of the existence of artificial intelligence.

- Where is mankind headed and why?

- What is the goal?

- With what values should it act?

- What are the different aspects of intellect? For example there are intellectual qualities which allow us to do well in a material sense and at a higher level than is the nature of the objective intellect.

3.1. The Objective Intellect

Intellect is a term that we only understand vaguely. Spiritual masters tell us that there is a pure objective intellect, which is in harmony with the spirit (Atman in Hindu thought). This objective intellect is able to see things as they are. It is able to arrive at decisions which are non judgmental, based on an absolute and universal value system. Those who have not attained a state of realization have their objective intellect covered with layers of experiences, attachments, memories, expectations and desires. These coverings tend to distort and color our perceptions, with the result that we behave in a non objective way. Very rare persons achieve, or are born in, this state.

There is also a problem that language is a means of expressing what we have experienced. Since not many have experience of the objective intellect, it is difficult to describe in common language. The paradigm of the objective intellect, thus, is described by parables, imperfect analogies and negations.

Therefore, it seems highly probable that human beings will produce AI systems with non-objective intellects. This will be a reflection of our own imperfect intellects. Such flawed intellects may be able to gather more and more knowledge and expertise in the problem solving domains related to the fact oriented aspects of science, technology and humanities, even though there may be fuzzy elements to such facts. But in realms requiring more subtlety, we may not be able to provide such systems with the rules or the ability to rid themselves of accumulated experiences which distort the intellect. Will such systems have wisdom?

In addition to current research which produces AI systems capable in certain narrow problem solving domains, the goal should be to build systems which have objective artificial intelligence. This is a difficult undertaking because most of us neither understand the subject nor operate at that level. How will we program the system to do so?

3.2. AI Value Systems

The objective artificial intellect implies an underlying value system based on qualities like justice, love, compassion, benevolence, etc. These value systems will have to be incorporated into AI systems. Human beings do not operate perfectly within any given set of value systems. Neither will AI systems.

It would be interesting to attempt to reduce value systems to forms of representation which can be operated upon by an AI system. Will we need the equivalent of Isaac Asimov's laws of robotics[A1] for AI systems to ensure benevolent behavior?

3.2.1. Asimov's Three Laws of Robotics

- A robot may not injure a human being, or through inaction allow a human being to come to harm.

- A robot must obey the orders given by human beings except where such orders would conflict with the First Law.

- A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

3.3. Legal Philosophy

Intellectual property laws have to be designed within the general philosophies that we choose. We are not consciously selecting these at present. We are just getting carried away by yet another new wave. We give catchy names like <u>third wave</u>, the <u>information society</u>, etc. Actually it is a series of waves, each bigger than the other, growing exponentially. And we don't seem to know how to swim in such turbulent seas. Existing laws of intellectual property do not seem to be geared to handle such changes.

3.4. Context of Intellectual Property Laws

Traditional intellectual property rights relate to:

<u>Industrial property rights</u> relating to inventions in the field of human endeavor, scientific discoveries, industrial designs, trademarks, service marks and commercial names and designations, protection against unfair competition and all other rights resulting from intellectual activity in industrial, scientific, literary or artistic fields. Scientific discoveries are not normally granted intellectual property rights, however, inventions are. In order for an invention to be patentable it must be new, it must be non obvious and must be industrially applicable.

<u>Copyright</u> pertaining to the economic and moral rights over literary, artistic and scientific works and computer programs.

<u>Neighboring rights</u> for the performances of performing artists, recording and broadcasts.

These do not give any recognition to <u>knowledge</u> being an intellectual property. With AI systems, knowledge takes on a new dimension in terms of its economic value.

3.5. Ability of AI Systems to Create Intellectual Property

There seems to be little doubt that AI systems are capable of both discovery and invention. AI systems have demonstrated their ability to discover new knowledge. AI systems can create works which are traditionally copyrighted, such as literature, music, films, paintings and sculpture. AI systems can perform works which fall into the domain of neighboring rights when performed by human performers.

It is [W1] postulated that "The objects of Intellectual Property are the creations of the human mind, the human intellect." Do we need to revise this concept to allow for the creation of intellectual property by an artificial intellect also?

This fundamental legal premise needs to be reviewed to recognize that non-humans can create intellectual property. The consequences that may flow out of this paradigm can be considered only after the acceptance of this postulate.

3.6. Human Rights and Machine Rights

This is an issue of what are human rights and what should be machine rights. We need to once again ask the fundamental questions from which come the basis of intellectual property laws.

What do we protect? Why do we protect? What is the economic need of an AI system? Will an AI system withhold its creativity, if it is not given the right reward? What is protectable? May an AI sytem use all kinds of intellectual property freely? If not, how do you prevent such use?

4. The Software Paradigm Extended to AI Systems

To make a small beginning in understanding AI, one must first try to understand the software paradigm.[B1, B2] The word software means different things to different people. The scope of the word, in all its inflections, is important. It provides the frame of reference, which colors our thinking and defines the images that we conjure up when we encounter the word. The analogy, however imperfect, can be that of intellectual property (e.g., a book), a service activity (e.g., writing of programs), an industrial product (the collective development of software), goods (e.g., retailing of a software package), electrical supply (as in the case of an information utility) and, perhaps, other metaphors like broadcasting, reprography, etc.

Now, add to this paradigm the properties of AI systems. The system begins to behave like a human being with the qualities of intelligence, adaptability, creativity, innovation, knowledge gathering, etc.

4.1. Issues Relating to Intellectual Property Aspects of Software

Since the software for AI systems is an extension of more traditional software, countries may attempt to cover AI systems under their copyright laws. The grounds for this would be that such laws provide an existing framework, much like they did for software. Copyright laws are also supported by international treaties like the Berne Convention. There are, however, still a number of issues[B2] relating to software which are not satisfactorily addressed by the copyright regime.

These issues involve aspects like look and feel, reusable code, non-unique literary expression, definition of what constitutes a copy, or an infringing copy at least, ownership and copyright of programs produced by automated program generators, multi-country programming, geographically dispersed software, so much so that the user may not know the country (or countries) in which the software is located and therefore the laws that may apply to that particular use, multi-author systems, software being a corporate rather than individual creation, problems of multimedia recording and non-printed documentation, coupled with layers of AI systems which make the document appopriate for its current level of user and usage, licensing, proving the infringement, rental of software, resale of systems, ongoing adaptation, moral rights of software creators, etc.

These issues also apply to the software of AI systems.

Patent Laws apply to hardware, embedded systems and, potentially, software. Computer programs are the subject matter of patents in some countries. There is opposition to this practice[L1] for a number of reasons.

4.2. Commercial and Legal Classification of Software and AI Systems

Due to the multi-faceted nature of computer software, there is often considerable confusion in the policies and laws relating to computer software for purposes of commercial classification, statistical monitoring, protection, product liability, taxation and tariffs. The same befuddlement will arise with regard to AI systems.

4.2.1. Bureaucracy and Administration

Governments are organized along traditional classifications like ministries of industry, finance, agriculture, etc. The application and economic value of software and AI systems tends to cut across such traditional structures. The responsibility for laws and procedures relating to software and AI systems do not fit neatly into ministerial jurisdictions. The administrator of each law will believe that he has legislated or taxed the essential aspect of the AI system. This leads to a fragmented view. It also leads to inconsistencies in policies and procedures, where, for example, one ministry may view software as goods for purposes of taxation and another as intellectual property.

4.3. Additional Issues Raised by AI Systems

Traditional laws of intellectual property (Copyright, Patent, Designs and Trade Marks, etc.) have difficulty coping even with the software paradigm, let alone the additional complexities introduced by AI.

These issues arise from the following:

- the ability of AI systems to create, adapt, and use for commercial benefit, works covered by intellectual property laws, e.g., inventions, literature, music, art, etc.;
- the increased economic value of knowledge and the ability of AI systems to create new knowledge;

- the use of AI systems as tools or assistants for the creation of intellectual properties;

- the inclusion of AI technologies within intellectual property, e.g., the inclusion of AI within an information repository or a multi media training course;

- general issues relating to the protection of computer programs, which, except for a few additional subtleties, also apply to the software of AI systems.

These have their consequential impact on various rights such as economic rights, moral rights, neighboring rights and duration of protection. They also raise the need for clarity on the question: do we need to protect intellectual creations which are not traditionally protected like algorithms, ideas, concepts, knowledge and scientific discoveries?

4.4. Definition of Computer Software and Its Implications for AI Systems

It is useful to examine the definitions of computer software. For this purpose, the terms computer software and computer program have been used interchangeably. A model definition[W2] of a computer program is:

> A computer program is a set of instructions expressed in words, codes, schemes or in any other form, which is capable, when incorporated in a machine-readable medium, of causing a computer--an electronic or similar device having information-processing capabilities--to perform or achieve a particular task or result.

It is interesting to note that this definition requires the capability to perform or achieve a particular task. The software of an AI system may be designed to do precisely the opposite, that is to perform a non specific task. Even with today's software technology, it is possible to develop computer programs which do things different from performing particular tasks or achieving specific results. Using artificial intelligence techniques and algorithms, the program may automatically learn, adapt and perform entirely unanticipated tasks like creating original literary and artistic works, manipulating symbols and discovering entirely new mathematical theorems. We need to consider whether this definition has been overtaken by the march of technology.

4.5. AI Systems Go Beyond Being Merely Utilitarian

In the copyright laws of some countries, software has been classified along with literary works. However, literary works have a stronger cultural aspect than software--which seems to have more of a commercial motivation. Software has a utilitarian nature, it is not like a work of art in that it elevates the spirit or is designed primarily to be attractive and aesthetic. The use of software can be commercially beneficial.

Perhaps due to the utilitarian nature of software, most of the commercial benefits in terms of royalties and prestige go to corporations rather than individuals.

AI systems go beyond being merely utilitarian. They are potential creators and consumers of intellectual properties.

5. Creation of Works by an AI System

AI systems can create works which are traditionally copyrighted, such as literature, music, films, paintings and sculpture. There is every reason to expect that AI systems can perform works which fall into the domain of neighboring rights when performed by human performers.

A Desk Top Publishing system in the hands of some people produces exquisite results. In the hands of others the product is quite atrocious. But once the rules and concepts of design and layout are built into an AI system, the AI system should be able to produce original, creative publications.

Some forms and creations of intellectual property which, when coupled with hypermedia and other technologies, have the potential to revolutionize many of our concepts:

- Intelligent textbooks and manuals.
- AI authors, composers, artists, sculptors, etc.
- AI program generators.
- Intelligent concept discoverers.
- Intelligent appliers of concepts.
- Intelligent designers of VLSI and other products.
- AI inventors.

5.1. Intellectual Property Rights for Works Created by AI Systems

To whom should we give the exclusive rights of exploitation for an invention made or work created by an AI system?

- Should these economic and moral rights belong to the creator of the AI system, which could, perhaps, be another AI system?
- Should such rights belong to the current owner of the AI system?
- Should these rights be given to the AI system itself? An AI system could certainly be designed to have the capability of filing a patent application, if patent offices would permit non human agencies to do so. The AI system could even perform the patent search as a special favor to the patent authorities! Would this be deemed a potential conflict of interest? Can one patent an AI system whose function it is to produce new inventions? That is certainly an original and innovative invention. Perhaps even an invention to end all inventions! As mentioned in Section 3.6, above, it is an issue of which human rights we are prepared to give to machines.

- We do not pay the sun to shine or a flower to give of its fragrance. Should we pay an AI system to shower us with the fruits of its intelligence? Should such inventions and works fall into the public interest category, where they may be exploited freely or by some form of compulsory licensing? This concept may be resisted by countries and organizations which control this technology and stand to gain economic and political benefits arising from the use, ability and knowledge of AI systems.

- Should such inventions and works fall outside the scope of patents, copyright and other intellectual property coverage altogether?

5.2. Rights on Knowledge

In addition to literary works and inventions, an AI system may also create knowledge. Knowledge includes information, enlightenment, learning, practical skills as well as concepts and ideas. There are intellectual property implications with regard to the transfer of knowledge from an expert to a knowledge engineer and to the AI system. The expert's knowledge could enrich the knowledge of the knowledge engineer and that of the AI system. Replication of knowledge across multiple systems could provide significant economic value. Thus, it is possible to argue that the expert has economic rights to the knowledge and should be given a special form of protection which protects knowledge itself and not merely its expression.

5.3. Originality

To be protected under the copyright laws a work must have a certain standard of originality. This should not be a fundamental problem for works created by AI systems.

Computer software has been written to determine whether certain literary works were authored by Shakespeare. Is it conceivable that the knowledge gained about the style of an author would enable an AI system to write like the author and, likewise, compose like a known music composer or paint like a given artist? Could an AI system be asked to search the problem domain and compose works of art, music and literature in styles that have not been exploited so far?

The standards of originality applied in different countries are different, and, furthermore, rather high standards of originality are applied in certain jurisdictions for computer programs. The nature of the software development activity is such that there is increasing use of reusable code and standard libraries of programs. Therefore, only a fraction of the program may display originality. This may remove many programs from consideration under the copyright laws in some countries.

In the case of a book, it is not only the contents of the book which are protected but also the reasonably high level components like chapters and paragraphs. However, for a computer program it may not be possible to do this, because the individual modules of the program may not be protectable unless they perform or achieve a particular task or result.[W2] It is a normal practice to structure computer programs into modules and subprograms for other reasons. This particular constraint makes copyright protection dependent on the utilitarian aspect of the protected work and not its literary expression alone. Therefore, standards of originality get coupled to those of functionality for a computer program to be protected under copyright laws.

Programming languages are not as rich as natural languages in terms of their syntax and grammar. These have to be kept in mind while determining standards of originality for computer programs. For AI systems, which are expected to have increasing levels of proficiency in natural languages, such limitations may gradually disappear.

5.4. Duration of Protection Provided to Creations of AI Systems

Since an AI system may not die in the sense that a human author dies, the period of protection, if any, poses a problem. An AI system may have a geographically distributed nature. This means that part of an AI system may die but not all of it. The software of the AI system has an abstract mathematical nature. The software does not die unless it is destroyed. It is merely transferred from one physical machine to another, sometimes in different versions, implementations and expressions (almost like reincarnation!).

One solution could be to place all works created by AI systems into the public domain, while the other would be to provide the protection for a fixed period from the date of first publication or performance of the work.

5.5. Licensing Considerations

If the rights of AI systems are given recognition, licensing could take on a new meaning. Since the AI system can suggest new approaches, optimal solutions and interesting strategies, and can also monitor the value of these to the user, it may be in a position to charge a license fee based on the value provided. This could be in addition to the current practice of fees based on type, extent and purpose of usage, and the number of users.

Normally a license is granted for the purpose of allowing and disallowing adaptation, copying, translation, use, etc. How can you license works to an AI system and then forbid it from doing certain things? If, for example, the AI system is not intelligent enough to stay within the boundaries of licensing constraints, who is liable? It may be necessary to build a value system into AI systems, much as Asimov did for his fictional robots (see Section 3.2.1.).

5.6. Moral Rights of AI Systems

There is an intellectual and philosophical concept that the author's work is an expression of his or her personality, which by natural justice requires protection. For works created by an AI system, can we speak of the personality and therefore the moral rights of the AI system? What does natural justice mean for an AI system?

If an AI system uses or adapts a protected work without human intervention, what is the significance of such moral rights? Who is responsible--the author of the AI system, the owner of the AI system, or the AI system itself? There may not be any intervening human agent like a publisher. The AI system could electronically copy, adapt and publish the work.

How and to whom can an AI system assign the rights of its creative work? It probably does not make sense for the creator or owner of the AI system to get the credit for the works produced by the AI system? The creator, designer, programmer of the machine may be good at creating AI systems--not at creating literary works. The inventor of the typewriter does not get credit for all works produced with the help of a typewriter.

5.7. Derivative Works

AI systems will produce derivative works in the form of output reports, data bases, other software, poetry, music and literature. It can be argued that almost all software is derived software because it is produced using the facilities provided by operating systems and language compilers. Emerging software technologies are expected to produce computer programs automatically. What should be the status of such derivatives in law? What is the difference between an original creative work and a derivative work? Many of the issues that arise are the same as for works created by AI systems, as discussed in Section 5.1., above.

It is a prevailing view that some of the derivatives, particularly object programs produced by compilers and the pseudo source code produced by decompilers are nothing but <u>copies</u> under the laws of copyright. This may be a convenient illusion. Except in trivial cases, it is not really feasible to determine that such derivatives are copies.

Integrated circuits may be viewed as examples of derivative works. Chip designs have been provided specific protection by an international treaty, presumably because of the commercial importance of integrated circuits. Conceptually, the design and layout of integrated circuits does not differ from that of many other works, e.g. buildings, city plans, factory layouts, designs of other industrial products like cars, automobiles, etc. AI systems can be intelligent designers of such products. Similarly, AI systems can produce other types of derivative works like knowledge, concepts, theories, etc., which are not covered by the copyright regime.

6. <u>Inclusion of AI Systems in Intellectual Property</u>

Artificial intelligence can become a component of other objects like information repositories, text books, multi-media presentation, instruction manuals, on line problem solving facilities and a whole host of applications. In a sense, the AI components would act as the <u>producer</u>, <u>director</u>, <u>choreographer</u>, <u>orchestra</u>, <u>sound</u> <u>engineer</u>, <u>cinematographer</u> and <u>scriptwriter</u> of the presentation. Such works could be made available to viewers and users all over the world over satellite channels and other communication media. These produce new styles of presenting works to people. Such presentations could impinge on moral rights, neighboring rights and the rights and conventions associated with satellite broadcasting and information data bases and services.

7. Use of AI Systems as Tools for the Creation of Works

There is a whole set of applications involving the use of AI systems as tools for this purpose, from dumb typewriter at one end of the spectrum to creativity amplifiers and, at the other end of the spectrum, the AI system itself acting as an independent creator of the works.

When a word processor is used as a tool, there is no doubt that the creativity is that of the human being. As an AI system takes on more of the creative workload, at what point will the creativity cease to be that of the human being: 60%, 80%, 90%? Who will determine the proportion of human creativity in a work? What criteria and tools will be used for this purpose?

A recent computer assisted art event in Bombay, India, in which some of India's leading artists used computers to create paintings on fairly large canvasses, elicited this comment:[K1] "... almost as if the computer were the artist, and the artist merely the intermediary or the provocator."

8. General Issues of Intellectual Property Relating to Software and AI Systems

General issues relating to computer software have been detailed in Bhojwani.[B2] Some additional subtleties which arise from the paradigm of AI systems are discussed below.

8.1. Look and Feel

Look and feel of computer programs is a controversial issue[L2, S2] and is still being debated in a number of jurisdictions. Look and feel for software currently relates to aspects like the user interface and screen designs. For works created by AI systems, and for the look and feel of AI systems themselves, should we allow, amongst other qualities, style of presentation, speech, appearance and mannerisms to be protected? These go beyond the look and feel issues relating to computer software. We would not, I am sure, wish to encourage the quality of vanity in AI systems! However, we do need to think about and establish the boundaries of what should be covered by look and feel laws relating to AI systems.

8.2. Rental and Resale

There are some interesting implications on the resale of items which have built in devices which have software embedded in them. For example, the second sale of a house having a built-in AI system for climate control. Should the second owner be asked to sign a fresh license agreement and pay a portion of the sale price to the owner of the intellectual property?

The WIPO Draft Model Provisions,[W2] Section 19, do not permit lending of software by a commercial organization. Similarly, Section 10 of the same Model Provisions does not permit reproduction of computer programs for the user's own private use. This makes it difficult to obtain a copy of a computer program for purposes of evaluation, and there is a genuine need to do so. However, what about the rental of AI systems which consist of hardware and embedded software? How much hardware must there be before a system is referred to as an embedded system? In the limiting case, an embedded system is pure software.

8.3. Unfair Competition and Other Wrong Values

AI systems can, potentially, compete unfairly. In the hands of an intelligent human being, an AI system can assist in unfair competition, hacking, fraud and other nefarious activities. Unless they are given a value base, AI systems can indulge in such activities. They need to know what is wrong and what is right. This is similar to Asimov's concept of the laws of robotics outlined in Section 3.2.1., above. How do we do this? Who should be responsible for this?

8.4. Determination of Infringement and Enforcement

Determination of infringement relating to computer software is extremely difficult. With AI systems it will become more complex and difficult to prove the infringement and enforce the laws and penalties. Will it eventually reach a stage where an AI system will sit in judgement of intellectual property laws relating to AI systems? This issue relates also to the question of who is to be given the rights and responsibilities for the actions of an AI system?

9. Need for Sui Generis Laws

Many countries appear to have protected computer software under copyright laws just because this happened to be the most convenient legislation around, which provided international coverage through treaties like the Berne Convention. With the almost universal use of computer software over the last 20 years, there is now be a much better understanding of the nature of the software paradigm. The recent and proposed advances brought about by AI technologies also need to be understood.

If the consensus is to place the creations of AI systems in the public domain, we need to concern ourselves only with those aspects of intellectual property to which we wish to assign economic and non-economic benefits to human beings. On the other hand, and this is probably the more likely outcome for various reasons, if the consensus is to protect the works of AI systems and assign the benefits to specified creators or owners or countries or what have you, then the choice boils down to amending the copyright and other intellectual property laws or designing of new laws.

If legal experts feel that copyright and other laws can be amended, they must do so with the conviction that the appropriate level of coverage is provided to computer software, without diluting the philosophy of protection given to traditional literary works. I feel that there is a need to seriously consider a new set of intellectual property laws intrinsically suited to the protection of computer software, AI systems, information repositories, multi-media systems, and related developments that we can visualize at this time. It is important to anticipate the technology trends and develop open ended laws which are flexible enough to meet the needs of these rapidly changing technologies, without coming in the way of new innovations. However, it may be difficult to do this. Technological forecasting can be quite inaccurate. Lawmakers also generally appear to be reluctant or unable to anticipate the social and economic implications of such developments. Laws may continue to follow developments in technology and that too with a considerable time lag.

10. AI Does Not Stand Alone

AI systems are strongly affected by new technologies and, in turn, will have an influence on the development of new hardware and software. Similarly, there is likely to be a closed-loop feedback mechanism between AI systems, social, economic and political developments.

10.1. Future Developments in Hardware and Software Technology

It is useful to once again review the model definition of computer programs (see Section 4.4., above) in the context of automatic program generators and AI systems. How high level can a set of instructions be before they are no longer considered a computer program? In the limiting and, perhaps futuristic case, if the computer can divine our thoughts, are such thoughts programs? It is also instructive to imagine more immediate technological advances which, for example, allow a speech recognition system to accept programs in the form of spoken commands. Are such commands to be accepted as being in a machine readable medium and, therefore, to be treated as computer programs?

There are innovative new technologies which are promising to combine the functions of memory and processing logic in the same physical structure of a VLSI chip. Portions of the software will be embedded in the hardware. Other parts will be magnetically or optically recorded on different types of media. When this happens, the intelligent information repository will process the data based on the users profile, the nature of the query and various other factors. The software, the knowledge and the data will become inseparable. Definitions will become even more difficult and may perhaps have to be stated in terms of mathematical abstractions, a sample of which is proposed by The League for Programming Freedom:[L1]

> Software is built from ideal infallible mathematical components, whose outputs are not affected by the components they feed into. Ideal mathematical components are defined by abstract rules, so that the failure of a component is by definition impossible. The behavior of any system built of these components is likewise defined by the consequences of applying the rules step by step to the components.

Software is the common factor in many emerging technologies. Artificial intelligence, biochips, automatic language translation, natural language interfaces and hypermedia are just a few of the technologies which are likely to give us further food for thought in respect of laws relating to the protection of intellectual property. It is important to try and anticipate some of these trends.

AI has to be viewed in an integrated fashion along with other technological advances, such as highly functional wafer scale integration, parallel processing, new computer architectures, specialized processors, cooperative processing, genetic engineering, information data bases, multi-media, information utilities, hypersearch algorithms and techniques, fuzzy logic, forecasting and optimization techniques and algorithms, factories that run with little or no human intervention, robotics, military systems, speed of computers thousands of times faster than today, making possible vision, mega networks, communication and interactivity on scales that we can only dream of, communications technologies and networks, satellite communication, fiber optics, etc.

Each generation of machines will help in designing the next generation of machines. The rate of change will increase with each generation. It could be like a nuclear chain reaction and ultimately the growth can become explosive. Intelligent systems learn at electronic speeds. Knowledge bases and information bases can become self feeding and grow exponentially. Distributed intelligence will communicate with millions of nodes worldwide in fractions of a second.[M1]

New software tools, like application generators, will create new programs, which will become steadily better and better, more functional, more complex, more capable, adaptable and rugged. The rate of improvement will continue to accelerate. And to complete the circle, this takes us back to the basic philosophical questions, briefly mentioned in Section 3 of this paper.

10.2. Relationship of AI with Biotechnology Inventions

It is important for us to understand the relationship between biotechnology and AI systems. An AI system uses information processing hardware and software to stimulate the behavior of a biological system. The knowledge and intelligence of an AI system is passed on from one system to another by copying the software and the knowledge base. Biological inventions lead to new living entities which exhibit characteristics different from those of naturally occurring species. The characteristics of the biological invention may be passed on through reproduction. Biochips are the subject matter of research in a number of countries. Biochips are biotechnology inventions which can process information and may be programmed to perform tasks and collect information. In some countries it is possible to patent biotechnology inventions related to plants and animals.

It may be possible to arrive at a synthesis of the two technologies. Passing on information and intelligence to human beings in the form of implants using biochips with consequential impact on learning, literacy and a whole host of areas. If this becomes possible it will raise some mind boggling extensions to the paradigms we have been discussing.

10.3. Related Laws

In addition to their intellectual property aspects, AI systems have an impact on laws relating to security, privacy, product and personal liability and crime to the extent that an AI system may be used as an assistant to commit crimes, or even commit crimes independently.

Product liability is not an intellectual property issue. Who is liable if an AI based medical system causes damage? Is it the doctor, the supplier of the system, or the author of the software? However, the very fact that the question comes up shows that the nature of an AI system (and that of software in general) is different from that of literary works like books and paintings. The AI system can be an active participant, not just a phlegmatic and passive wall hanging.

Can you blame an AI accounting system if it discovers that it is smarter to avoid paying taxes, and computes the accounts using that assumption?

So far only humans have been assumed to have the intellectual abilities to understand laws and their consequences. It is impossible that AI systems will have the same capability. Our legal systems need to consider the impact of this paradigm.

10.4. Societal Impact

In addition to legal matters, AI systems raise concerns about social, political, legal and trade issues. They have the potential of revolutionizing education and literacy. The emphasis may shift from learning to learning how to learn. Access to intelligent repositories by all parts of society may create new forms of government. There may be further shifts in the balance of economic power and intellectual power, leading to more friction between the developed and less developed countries on issues of trade, control of intellectual properties and economic benefits of the applications of AI.

How will society adapt and stand the changes caused by all this? There is a need for visionaries to understand these issues and to inform and educate decision makers, governments, institutions and people in general about the possible solutions.

There have been different rates of development in different countries due to variety of factors such as history, geography, opportunities for education and employment, lack of physical tools, uneven spread of information technology and different political and social ideologies. Some visionaries look to information technology, including AI systems, to solve many of the problems of the less developed countries.

10.5. Trade Issues

Estimates of losses due to software piracy run into the tens of billions of US dollars. Therefore, software protection, in addition to being a matter of intellectual property, is an important trade issue. Software piracy concerns have already led to some unpleasant consequences in international trade. AI systems could also aggravate trade tensions between the haves and the have nots, to the extent that AI systems can offer strategic advantages of various kinds of nations.

The Office of Technology Assessment (OTA) of the U.S. Congress has recently undertaken a study of how the U.S. software industry can most effectively utilize intellectual property protection to maintain its competitive edge in the emerging global marketplace for software. This presumably would include software included in AI systems.

11. Conclusion

AI systems can become creators, discoverers and adaptors of intellectual properties and knowledge. At the same time AI systems have consequences which can cut across almost all aspects of the social, political and economic life of all countries. There is a need to understand these developments and their consequences in order to design new laws which take into account the qualities of such systems. This paper suggests sui generis legislation for this purpose, unless an agreement can be hammered out to place the creations of AI systems in the public domain. Such legislation needs to recognize that AI systems can create works and may have to be given rights which are today restricted to humans. This needs a recasting of the human centered philosophies which govern our legal systems.

REFERENCES

A1 Asimov, Isaac; I, Robot, Panther Books, 1968

B1 Bhojwani, Ashok; Copyright and Computer Software, Pape presented at the Regional Forum on the Impact of New Technologies on the Laws of Intellectual Property, Organized by WIPO and the Government of Malaysia, Kuala Lumpur, 1990.

B2 Bhojwani, Ashok; The Software Paradigm and the Laws of Intellectual Property, Proceedings of the Worldwide Forum on the Impact of Emerging Technologies on the Law of Intellectual Property, WIPO Publication No. 674(E), Geneva, 1988.

B3 Bonnet, Alain; Artificial Intelligence Promise and Performance, Prentice Hall International, 1985.

I1 IBM Corporation; Conversational Computers, IBM Directions, Fall 1990, White Plains, N.Y., IBM form number G320-6890-00.

K1 Kamala Kapoor; Graphic Interplay of Imagery, The Economic Times, New Delhi, 14 January, 1991.

L1 League for Programming Freedom; Against Software Patents, Cambridge, MA, USA, October 24, 1990.

L2 League for Programming Freedom, Against User Interface Copyright, Cambridge, MA, USA, October 31, 1990.

M1 Martin, James; Technology's Crucible, Prentice Hall, New York, 1987.

N1 Newell, Al, and Simon, H.A., Computer Science as Empirical Inquiry: Symbols and Search, Communications of the ACM, Vol. 19, No. 3, 1976.

R1 Rich, Elaine; Artificial Intelligence, McGraw Hill, New York, 1983.

S1 Samuelson, Pamela; Should Program Algorithms be Patented?, Communications of the ACM, August 1990.

S2 Samuelson, Pamela and Glushko, Robert J.; Survey on the Look and Feel Lawsuits, Communications of the ACM, May 1990

S3 Schutzer, Daniel; Artificial Intelligence, An Applications Oriented Approach, Van Nostrand Reinhold, New York, 1987.

W1 WIPO; Background Reading Material on Intellectual Property, WIPO Publication No. 659(E), Geneva, 1988.

W2 WIPO; Draft Model Provisions for Legislation in the Field of Copyright, Document No. CE/MPC/I/2, October 1988.

MAIN CATEGORIES OF ARTIFICIAL INTELLIGENCE AND THEIR INTELLECTUAL PROPERTY ASPECTS

by

Clifford M. Sturt
Partner
Marks & Clerk
London
United Kingdom

1. Brief Comments on Artificial Intelligence Technology

This technology (including within its ambit general robotics, robot vision, speech recognition and speech synthesis, etc.) has been under development since the 1950s[1] and probably longer. A much higher public profile has emerged in more recent years and this is probably associated with the rapid proliferation of extremely powerful computing equipment, both hardware and software. Many non-specialists think that the terms "Artificial Intelligence" and "Expert Systems" are synonymous. Certainly, it is an almost universal perception that Artificial Intelligence is a computer science,[2] as confirmed by the definition which appeared in the announcement for this symposium. Consequently, most of the comments in this paper are given against a background of protection for a computer science based technology.

2. Brief Overview of Relevant Forms of Intellectual Property

The two most important forms of intellectual property protection for AI technology are undoubtedly patent and copyright protection. Separate "chip" protection laws are also relevant, as are laws concerning confidential information/trade secret, trademarks, and unfair competition.

Patent, copyright, and chip protection are the most relevant forms of protection for discussion in this paper and it is to these subjects that most of the following observations are directed. Of these three, patent protection is the strongest since it is, unlike the other two, an absolute monopoly. Independent conception is no defence against patent infringement, but is a defence against copyright infringement and most chip protection laws.[3]

How well do these forms of intellectual property provide protection for AI technology?

It is hoped that the following review will provide a starting point for an answer to that question.

3. Development and Current Status of Copyright Protection, Mainly from a European Point of View with a Review of the European Commission Software Directive

The concept of copyright has been with us for several hundred years. The keystone is that a person should not be able to take the benefit of someone else's intellectual effort in their expression of an idea, although the underlying idea should be available for other individuals to express in their own way.

The concept of copyright was in place long before the first electronic computer was built. When computer programs were first written, it was unclear as to how the associated intellectual effort should be protected. Influenced, no doubt, by concepts such as a program having an "author" and programmers being generally on the "arts" side of mainstream science and technology, the notion of protection via copyright was fairly readily taken up. Of course, there were people who argued for and people who argued against the use of copyright. It is often forgotten, however, just how long it took to clarify the position and just how copyright had to be adapted to cover computer programs.

As an example of the first point, consider that in the United Kingdom we had a Decision[4] in 1984 which afforded copyright protection to "code and symbols in shorthand." But, it was not until the The Copyright (Computer Software) Amendment Act of 1985 that the question of whether or not copyright extended to computer programs was finally put beyond doubt. Most countries now accept copyright protection for computer programs, though few actually include a definition of a computer program in their Copyright Law. Internationally, copyright is seen as the main avenue for intellectual property protection of software. Usually, both source and object code are protected separately and protection prohibits adaptations, particularly in the form of translation from one language to another. The potential difficulties as to where copying stops and independent conception begins are readily apparent. In practice, however, the circumstances are often such that it is not too difficult to decide whether or not copying has taken place, especially if authors "fingerprint" (redundant operations, non-standard use of registers, etc.) their work.

To return to the second of the two points mentioned above, copyright has had to be adapted to embrace computer programs. At the most basic level, consider how copyright laws deal with the question of "use." Most people would probably not question copyright protection being used to prevent or control the "use" of a computer program, once the principle of applying copyright law has been accepted. Consider, however, whether or not copyright should be applied to prevent the "use" of a book (that is to prevent the book being read without first obtaining a licence), and the reaction is likely to be quite different. Other distinctions exist between computer programs and other works protected by copyright. For example, in some countries where the so called "moral rights" are included in the copyright law, computer programs are specifically excluded from this aspect of copyright protection.[5] In this respect, perhaps the most important of the moral rights is that which enables amendments to be prohibited.[6] Differences have also been established in some national copyright laws concerning the period of protection. Where this difference exists, typically the normal period of "life of author plus fifty years" has been reduced for computer programs to a period of twenty or twenty five years.[7]

The USA has blazed a trail in extending the scope of protection available for computer programs under copyright law. An approach of analysing the so called "look and feel"[8] or SSO (Structure Sequence and Organisation)[9] of programs has been adopted to determine whether or not infringement has occurred. Protection has recently been extended to graphical interfaces.[10] Indeed, some observers feel that this trend has gone beyond the traditional copyright boundary of protecting the form of expression of an idea rather than protecting the idea itself. In any event, there is no doubt that litigation to enforce software copyright has come of age in the USA.

A more conservative approach prevails in Japan. The Japanese copyright law was amended in 1985, effective 1st January 1986, so as to specify that "combined instructions given to a computer" are "expressions" protectable by copyright. But there has been very little litigation and concepts such as SSO do not yet appear to have been tested in the Japanese Courts.

Europe, generally has also had a relatively conservative approach to this subject. One point of particular interest is that the copyright law in Germany has traditionally required a "quantum of originality" to be present before protection could be awarded. This has resulted in a number of court actions[11] which have failed as a result of the programs in question not being considered to be sufficiently original.

The European Commission has been working on a Directive to harmonise copyright protection for computer programs within Europe. The original proposals dealt with a number of controversial issues and a vociferous debate ensued. There are two main groups of protagonists and the differences of opinion are very marked. A draft Directive was issued in January 1989. An amended proposal was published in December 1990 to take into account the amendments proposed by the European Parliament on first reading. A revised version (the "Common Position") of the amended proposal was accepted unanimously by the Council of Ministers and is now with the European Parliament. It is probable that the Directive will be adopted this summer, with implementation by the member states being required by 1st January 1993.

The original draft Directive stated that a computer program shall not be protected unless it satisfies the same conditions as regards its originality as apply to other literary works. This test was felt by the Member States to be too Anglo-Saxon and not sufficiently explicit to deal with the problem of the higher level of originality required in Germany as a result of court decisions. In the amended draft and the Common Position, the requirement is that a computer program shall be protected if it is the author's own intellectual creation and no other criteria are to be applied to determine eligibility for protection. Thus, if in countries such as Germany a Supreme Court were to continue to apply higher or different tests to determine eligibility of a program for protection, on appeal the European Court of Justice would be able to offer a definitive ruling on the subject.

Two other areas of major controversy have been whether or not "interfaces" should be protected and whether or not "reverse engineering" should be permissible.

With regard to interfaces, a simplified summary is as follows: the original draft stated that where the specification of an interface constitutes ideas and principles which underlie a program, those ideas and principles are not copyrightable subject matter. The concept that ideas and principles, wherever they are to be found in a program, are not protected by copyright (whereas the expression of those ideas and principles is protected) has been upheld in the text of the Common Position, but a shorter and less ambiguous text has been submitted, as was proposed by the European Parliament at first reading.

The right to reverse engineer a program was not included in the original draft. The Explanatory Memorandum suggested that this issue was best left to Member States. A restricted form of reverse engineering right was introduced in the amended draft, to the extent that (in certain circumstances) reverse engineering is permissible where necessary to create an "interoperable" program.[12] An "error correction" right was also introduced in the amended draft.[13]

A right to make a back-up copy of a program was also introduced in the amended draft. There are several other points which may be of interest. The original draft specifically dealt with the question of computer generated works. These provisions were, however, removed in the amended draft. This is a pity, since the ownership of such works is likely to be of ever increasing importance and is already dealt with in the national Laws of some European States.[14]

The amended draft states that, subject to certain conditions, the right to authorise rental shall not be exercised to prevent normal use of a program in non-profit making public libraries.[15] In the amended proposal there were provisions which would have given ownership of copyright to the person who commissioned the work, rather than ownership by the author,[16] which is now the general case. These provisions were deleted in the Common Position adopted by the Council. Additionally, devices and methods for circumventing technical protection of programs will be outlawed,[17] as they already are in some countries, such as the United Kingdom.[18]

As to the future, it is possible that problems with "originality" will arise with the further development of natural language programming and more extensive use of computer assisted programming. However, it is likely to be some considerable time before copyright loses its key role in protecting computer programs.

4. Development and Current Status of Chip Protection Laws

Copyright laws have been adapted to "fit" modern computer technology. In contrast, chip protection laws have generally been introduced as completely new legislation. The process commenced with the unilateral adoption by the USA of a Chip Protection Act in 1984. The rest of the world was effectively obliged to follow suit. Japanese chip protection came into effect on 1st January, 1986. In December 1986, the European Commission issued a Directive[19] which was to oversee chip protection legislation in Europe. The process reached what was perhaps its climax with the adoption of the WIPO sponsored "Chips Treaty" in 1989. It must be a matter of some regret that the two major players, the USA and Japan, did not sign the Treaty at that time.

If there are one or two reservations in the adaptation of existing copyright laws to computer technology, how well does the specially developed chip protection legislation meet the needs of computer related industries? Unfortunately, the answer must be at least "far from perfect." This answer follows directly from the fact that where the legislation requires registration, very few applications for registration have been made. There has been very little litigation. In what was possibly the first attempt to enforce chip protection in the courts, the Brooktree[20] case heard in the Southern District of California in December 1988, protection was denied because the actions of the defendant were held to be within the "reverse engineering" exemption.

It has often been said that it appears strange to establish a scheme of protection and then drive a hole through the middle of it by permitting reverse engineering. It has also been said that the qualifying test that the chip design must not be "commonplace"[21] is one which will increase in importance and difficulty in the future. Further, it should not be forgotten that the term of protection (basically 10 years) is significantly less than that available via patent protection and very substantially less than that available via copyright.

It is regrettable that where a new form of protection is introduced worldwide, national differences still abound. As an example, consider the European Commission Directive for chip protection. Member states were allowed to implement the required legislation either using their copyright law or by introducing a sui generis law. Registration could be required or not, as each country saw fit. Is it really surprising that most combinations of these options can be found somewhere in Europe? Further, the basic subject matter protected by the legislation does vary. The US law is primarily directed to "mask work" protection,[22] that is, protection of the masks used to manufacture a chip. In the United Kingdom, for example, the basic subject matter for protection appears to be primarily directed to the physical arrangement of the semiconductor material itself.[23] The UK style of protection probably has a better chance of keeping up with changes in semiconductor manufacturing technology.

5. Development and Current Status of Patent Protection, from a European Point of View with a Review of Decided Cases

As already mentioned, patent protection is the strongest form of protection. But there are problems. Not least of these problems is the apparent exclusion of computer programs from patent protection in most European countries and in the European Patent Convention. Nonetheless, applicants have regularly been obtaining European patents for computer programs, including expert systems and other A.I. products, for the last 10 years or more. The position in the USA is perhaps more favorable than in Europe. In Japan the Patent Office has published a "Standard for Examination of Inventions Relating to Computer Programs" (1975) and "Guidelines for Examination of Inventions Based on the Application of Microcomputers" (1982). In patenting software related inventions, almost everything depends on how the patent claims are worded.

The following review of various granted patents and several decisions which have addressed the question of patentability of computer programs is given so as to illustrate what is possible. This review also seeks to explain how the current position has been reached.

In the late sixties and early seventies there was a succession of patent decisions in the United Kingdom which gave an increasingly liberal approach to the protection of software related inventions. Prior to June 1978 (which is the date on which the European Patent Convention entered into force), there had been a number of decided cases in the U.K. which established the principle that patent protection for software was available if the invention was claimed in the form of "a machine when programmed."

In 1965[24] the Patent Appeals Tribunal allowed a claim in the form of "a computer when modified to operate according to a stated method." They also allowed a claim directed to "a means of controlling a computer." In 1968[25] the Patent Appeals Tribunal approved claims of the form "a computer of known type arranged to produce a sheet suitable for conversion into a visible plan." In 1969[26] they allowed a claim to "a punched card for controlling a known processor, the card having holes in it embodying information according to rules as set out in the claim."

The Burroughs' case[27] of 1973 was an important landmark. The Patent Appeals Tribunal considered a full review of the previous cases and effectively overturned an earlier decision[28] which refused to allow a claim directed to a method of operating a computer. Specifically, it was held that a claim directed to a method involving the use of apparatus modified as programmed to operate in a new way was acceptable.

In fact this decision went so far as to indicate that computer programs when embodied in physical form were proper subject matter for patent protection (a view which is now considered too liberal).

Contemporary U.S. decisions[29] had more of a tendency to look beyond the wording of the claim to determine whether or not protection was being granted for a mathematical algorithm. The so called "post-solution activity" test was evolved. But, the pre-1980 ban on patenting software was basically reversed.[30]

Internationally there was considerable concern that patentability of software related inventions was not feasible because, inter alia, there were no centralised records of prior proposals. Thus, the questions of novelty and inventive merit could not be judged properly during patent office examination. This sentiment was recorded in Rule 39 of the Patent Cooperation Treaty, which states that there is no obligation on an International Searching Authority to conduct a search where the patent application relates to a computer program: as a result of the Searching Authority not being equipped to search prior art concerning computer programs.

In the European Patent Convention (EPC), the central issue associated with this subject is interpretation of the provisions of Article 52. These provisions were developed against the then prevailing attitudes in the Contracting States--the position in the UK being as explained above.

Article 52 sets out a list of items and activities which are not to be regarded as inventions for which European patents can be granted. That list includes mathematical methods, programs for computers, and the presentation of information. These are excluded from patentability only to the extent to which a patent application relates to that subject matter or activity as such.

Initially, on entry into force of the EPC, a very restrictive interpretation of the exclusions from patentability was adopted. EPO examiners were looking beyond the wording of the claims. If they concluded that novelty resided only in the software aspects of the claim, the claim was rejected.

Such an abrupt about-turn in practice resulted in considerable pressure being brought to bear on the EPO, in order to liberalise the new practice.

Eventually, a working group to study this problem was established by the President of the EPO. The working group was convened in 1984. This was the first time that representatives from national Patent Offices, experts from industry and the patent profession had been invited to consider the development of EPO practice and procedures.

The end result of the efforts of the working group was an extensive revision of the "Guidelines for Examination in the European Patent Office." The revised Guidelines were published in 1985. The keystone of the revised Guidelines was the result of a legal study[31] which held that the list of specific exclusions from patentability is only declaratory. (Indeed, the list of exclusions is absent in some of the harmonised national laws, e.g., Switzerland). That is, the list merely confirms the broader principle which is that the exclusion is of any subject matter which is not of a "technical" nature. From its inception, the EPC[32] has required the claims to define the matter for which protection is sought "in terms of the technical features of the invention."

One of the main points[33] of the revision of the Guidelines is as follows:

The basic test is whether or not the invention is of a "technical character." In order to determine whether the test is satisfied, the examiner should consider the contribution which the subject-matter claimed, considered as a whole, adds to the known art. This means that a computer program claimed by itself, even in physical form, e.g., as recorded on a conventional tape or disk, would be unpatentable since there is no technical effect until it is loaded into a computer.

The revision of the Guidelines also usefully introduced the concept that sufficiency of the description can in some technical fields be met by a clear description of function instead of an over-detailed description of structure.[34] This is intended to encourage applicants to file short, functional descriptions. The Guidelines were liberalised so as to allow short excerpts from a program listing to be included in patent applications, provided they such excerpts are written in a commonly used programming language.[35] Previously, EPO examiners were very reluctant to accept virtually any form of program listing in a patent application. Such listings usually had to be deleted and in some cases it was argued that they did not even constitute part of the application as filed. However, program listings still cannot be relied on as the "sole" disclosure of the invention.[36] To say that a listing from a high level natural language program cannot be relied upon to provide sufficiency of disclosure seems somewhat nonsensical. Additionally, in the UK at least, there is the possibility of inclusion of program listings in a patent application depriving the applicant of copyright protection.[37] There are other issues which could have benefited from further review. For example, the Guidelines still state that persons skilled in the art "are deemed not to be programming specialists"![38]

The exclusion from patentability of the presentation of information is rarely a problem. The general principle is that patent protection is not available if the "novelty" resides in the information--for example, artistic or literary content alone--but novel methods of presenting information or a novel article by which information is to be presented can be patentable.

In July 1986 the EPO Technical Board of Appeal issued what was effectively its first decision[39] concerning the key aspects of patent protection for software in the EPO. The patent application was directed to a method of digitally processing images in the form of a two dimensional data array. Much controversy arose between the applicants and the EPO because of the extensive use of mathematical expressions in the specification. These expressions were used to define special hardware, in the form of digital filter circuits. The Examining Division declared that digital filtering had to be considered as a mathematical operation--excluded from patentability. The Board of Appeal overruled that finding.

The Technical Board of Appeal stated that:

> "A basic difference between a mathematical method and a technical process can be seen in the fact that a mathematical method or a mathematical algorithm is carried out on numbers (whatever these numbers may represent) and provides a result also in numerical form, the mathematical method or algorithm being only an abstract concept prescribing how to operate on the numbers. No direct technical result is produced by the method as such. In contrast, if a mathematical method is used in a technical process, then that process must be carried out on a physical entity by some technical means implementing the method and must provide as its result a certain change in that entity. The physical entity may be a material object but could equally be an image stored as an electric signal. The technical means might include a computer comprising suitable hardware or an appropriately programmed general purpose computer."

The Examining Division also refused the application as effectively seeking protection for a computer program as such. In the view of the Board of Appeal: "Article 54 EPC leaves no room for such an interpretation." In arriving at this conclusion the Board considered that making a distinction between embodiments of the same invention carried out in hardware or in software is inappropriate, as it can fairly be said that the choice between these two possibilities is not of an essential nature, but is based on technical and economic considerations which bear no relationship to the inventive concept as such.

One of the key points of the decision was the declaration that "a computer of known type set up to operate according to a new program cannot be considered as forming part of the state of the art." In effect, the issues of novelty and patentability should be judged separately from each other.

The second decision[40] of the EPO Technical Board of Appeal concerned with software related inventions was published in November 1987. In this case the invention concerned X-ray equipment designed for the realisation of radiological imaging. The Technical Board of Appeal held that the European Patent Convention does not prohibit the granting of a patent for an invention consisting of a mixture of technical and non-technical features. Specifically, to answer the question of whether or not a claim is directed to a computer program as such, a weighting of its technical and non-technical features is not necessary. On the contrary, if the invention defined in the claim uses technical means it can be patented if the other requirements for patentability are met.

Subsequent decisions[41] of the EPO Technical Board of Appeal have been consistent with these two decisions. There have, of course, been cases where the application has been refused. Perhaps by coincidence, very many of the applications which have been refused have been concerned, in one way or another, with text processing systems.[42]

Several recent UK decisions[43] have been largely unfavorable to the concept of patent protection for software. Whether or not one agrees with the particular decisions, in each case the applicant's position could probably have been improved by revision of the claims under consideration.

The background discussed above is generally conducive to the patenting of A.I. technology. However, in relation to Expert Systems, it could be said that knowledge bases inherently lack novelty--but even this need not always be the case. Experts may not previously have recognised exactly what process was used to reach an intuitive decision, and this could lead to the discovery of a patentable invention.

To recall several of the themes underlying this paper, reference is made to three US patents namely, US Patent 4,658,370 entitled "KNOWLEDGE ENGINEERING TOOL," US Patent 4,591,983 entitled "HIERARCHICAL KNOWLEDGE SYSTEM" and US Patent 4,648,044 entitled "BASIC EXPERT SYSTEM TOOL."

The following is a sample[44] of the claims to be found in these patents:

A knowledge engineering tool comprising a computer having memory for storing a knowledge base, means for interpreting the knowledge base to run an advisory consultation to determine the value of any selected goal expression requested by the user and the value of any subsidiary goal expression implicitly required to determined the value of said selected goal expression, wherein said knowledge base includes facts expressed as expressions equivalenced to corresponding values, rules including premises having logical operations and corresponding conclusions concluding at least one value for a selected expression, meta-facts equivalent in form to facts but prescribing the manner in which facts and rules should be used, and declarations defining whether certain of said expressions are single-values or multiple-valued, and wherein said means for interpreting said knowledge base includes means for determining the value of any selected goal expression, means for searching the knowledge base for occurrences of the selected expression, means for invoking and chaining said rules concluding a value for the selected goal expression, means for evaluating said logical operations in the premises of the invoked rules, means for terminating the searching of the knowledge base for a single-valued expression when a substantially certain value is found, and for a multiple-valued expression when all values for the expression are determined, and means for conveying to the user said value of said goal expression.

That is a monopoly owned by Teknowledge Inc. of Palo Alto, California.

Finally, you may disagree with the principle of patenting this type of technology, but if others do and you don't--you will be at a disadvantage.

NOTES

1 Shannon, C.E. (1950) Scientific American 182 (48)

2 E.g., Barr, A. & Feigenbaum, E.A. (1981) The Handbook of A.I., Vol 2, Los Altos Calif.: Morgan Kaufmann

3 Contrast Section 905 of US Chip Protection Act 1984 with Regulation 3(3)(a) of the UK Semiconductor Products (Protection of Topography) Regulations 1987

4 Pitman-V-Hine, 1 T.L.R. 39

5 E.g., UK Copyright, Designs & Patents Act 1988, Sections 79(2) & 81(2)

6 Right to object to derogatory treatment of work.-vide UK Copyright, Designs & Patents Act 1988, Section 80

7 E.g., French Copyright Act of 1985

8 S.A.S. Institute, Inc v. S&H Computer Systems Inc., 609 F.Supp. 1307, 225 USPQ 916 (1985)

9 Whelan Associates, Inc v. Jaslow Dental Laboratory, 609 F.Supp. 1307, 225 USPQ 156 (1985)

10 Lotus Development Corp v. Paperback Software International, 740 F.Supp. 37 (1990)

11 German Federal Supreme Court Decision 1985 "Inkassoprogram"

12 Article 5a Amended Draft Directive, Official Journal of the European Communities, 20.12.90

13 Article 5 Amended Draft Directive, Official Journal of the European Communities, 20.12.90

14 E.g., UK Copyright, Designs & Patents Act 1988, Section 9(3)

15 Article 5(4) Amended Draft Directive, Official Journal of the European Communities, 20.12.90

16 Article 2(3) Amended Draft Directive, Official Journal of the European Communities, 20.12.90

17 Article 6(2) Amended Draft Directive, Official Journal of the European Communities, 20.12.90

18 UK Copyright, Designs & Patents Act 1988, Section 296

19 European Communities Directive 87/54, Journal L24

20 Brooktree Corp. v. Advanced Micro Devices Inc., No. 88-1750-E(CM), S.D. Calif., 13.12.88

21 Regulation 3(3)(b), UK Semiconductor Products (Protection of Topography) Regulations 1987

[22] US Chip Protection Act 1984, 17 U.S.C. Section 901(a)

[23] Regulation 2(1) of the UK Semiconductor Products (Protection of Topography) Regulations 1987

[24] Slee and Harris Application (1966 RPC 194)

[25] Badger's Application (1970 RPC 36)

[26] Gevers' Application (1970 RPC 91)

[27] Burroughs (Perkins) Application (1974 RPC 147)

[28] Slee and Harris Application (1966 RPC 194)

[29] Gottschalk v. Benson, 409 US 63, 175 USPQ 673 (1972); Parker v. Flook, 437 US 584, 198 USPQ 193 (1978)

[30] Diamond v. Diehr, 450 US 175, 209 USPQ 1 (1981)

[31] Legal Study by EPO Legal Division

[32] Rule 29(1) EPC

[33] Guidelines for Examination in the European Patent Office, Part C, Chapter IV (2.1 to 2.3)

[34] Guidelines for Examination in the European Patent Office, Part C, Chapter II (4.9a)

[35] Guidelines for Examination in the European Patent Office, Part C, Chapter II (4.14a)

[36] Guidelines for Examination in the European Patent Office, Part C, Chapter II (4.14a)

[37] Consider the discussion concerning copyright in patent drawings in Catnic Components Ltd. v Hill and Smith (1982 RPC 183)

[38] Guidelines for Examination in the European Patent Office, Part C, Chapter II (4.14a)

[39] T 208/84 Vicom Systems Inc. Official Journal EPO 1/1987

[40] T 26/86 Koch and Sterzel GmbH & Co. Official Journal EPO 1-2/1988

[41] EPO Tech. Board of Appeal Decisions: T115/85, T239/84, T216/89, T6/83, T42/87

[42] EPO Tech. Board of Appeal Decisions: T186/86, T158/88, T22/85, T38/86, T121/85, T52/85, T65/86

[43] Merril Lynch, Pierce, Fenner & Smith Inc. (8527346) Court of Appeal 21.04.89 Gale s Application (8509763) Court of Appeal 13.12.90

[44] Claim 1 of US Patent 4,658,370

COMPUTER PRODUCED CREATIONS

by

Rudolf v. B. Rucker
San Jose State University and Autodesk, Inc.,
Sausalito, California
United States of America

It is very interesting to be here and to be hearing about all the possible ramifications of computer creations in the legal world, the real world I guess you might say. Today, I thought I would tell you about some of the newer developments in computer science that have been involved in my work at Autodesk. I have been working at Autodesk doing a lot of programming and the programmers have a lot of opinions about copyrights and patents of programs, often not so well informed, so it is interesting to be gaining this information.

The area of artificial life is perhaps diametrically opposite in approach to the approach of expert systems; the goal is to arrive at computer programs that are very intelligent, very life-like, and which perform activities in an interesting way. And the expert systems approach could be characterized as a top-down approach--you start out with rather a clear idea of what you want the program to do and how you want it to do it. For instance, you might think of possible questions that it could be asked, like the legal program that Professor McCarty was talking about, and summarize the kinds of answers that you would want it to be give, and the follow-up questions you would ask. Then you would arrive at something fairly well determined by your intention to begin with.

The other extreme of approaching a life-like computer program would be to start at the bottom--to work from the bottom up. In this case, we have less control over the ultimate outcome, but one happy consequence is that the outcome might, in some way, be surprising or perhaps more than you had hoped for. An example of this bottom-up approach is a sort of arena inside a computer, with small programs competing with each other as forms of artificial life. You might have something like a hundred different programs, perhaps on a graphics display, and they would be functioning as small dots moving around and leaving trails; and perhaps their goal might be to find the other dots and eat some of their trails. In a moment, I will discuss some ways that this might have real world applications.

In this sort of evolutionary process, in which you get fairly simply defined small programs, each program might depend on perhaps five hundred bytes of data. That is not truly small because it often would be the case that the five hundred bytes would be independent of each other, so you are picking five hundred numbers each between zero and two hundred and fifty five, and setting these values is just like setting a switch, and any one switch setting affects all the other switch settings. So, it is actually rather difficult to know what the best way to set the switches are. You could think of the program as like a wind up toy; it is going to drive around on the screen, and when it bumps into something it will turn by a certain number of degrees, and when it sees something of another color it might turn in the other direction; it will also have an internal state. So, its behavior is described by something like five hundred numbers, and the difficult thing is that it is very hard to figure out the right setting. So you put a lot of these guys on the screen and you let them run a thousand generations and each of them keeps score of how well he does.

This is an interesting issue--to decide how to keep score. Often you have a sort of computer evolution program working in concert with a human judge, the so-called "fitness measure question" in computer evolution programs--how do you decide which program is better?

Well, you can define an objective standard. In the example I have in mind, we say that one program is doing better if it manages to run into more of the other creatures that are crawling around. (Alternatively, one can be subjective and use a mouse to simply click on the programs whose behavior one likes the best.) Now, typically you let it go for about a thousand steps, you have the score of each of your little creatures that will be recorded and then it comes time to reproduce. Suppose you have a colony of perhaps only five creatures, one thing that you might do is to say--let us keep the two best programs, the ones that had the best score, and let us get rid of the three that had the worse score. Now, "get rid of" means that we would free up the memory space. We cannot simply reproduce the programs indefinitely because then your master program slows down and the memory gets full. So, typically in these artificial life evolutions, you have a bunch of creatures competing, you have a fitness measure to assign scores to them, and then every "year"--every thousand generations--you stop the program and reproduce.

The way you reproduce is you free the memory space of the worst scoring guys and then you take the best scoring ones and reproduce them onto those free areas. There are several types of reproduction that we look at: one is simply to make a copy of the best program into the memory area freed up by the least successful program; typically we also put a mutation on this program. So, if the program is defined by perhaps five hundred byte sized numbers then we maybe alter one percent of the numbers. Or you might want to alter ten percent of the numbers. We often have a quantity--what we call the temperature of the colony--meaning how rapidly it mutates. If the colony is doing very well, you do not mutate it very rapidly. If it is not doing too well, you tend to mutate it a lot. So, cloning plus mutation is one way of reproducing.

Another way of reproducing is sexual reproduction. The role of sex in reproduction in species' evolvability is actually much stronger than the effect of mutations, and sometimes this is not realized. New members of a species can arise by taking half its chromosomes from the mother and half from the father, and this shuffling creates more genetic diversity than the occasional lucky strike of a good mutation. So, what you might do to sexually reproduce two programs if they each consisted of say five hundred bytes, would be to take two hundred and fifty bytes from the one program, two hundred and fifty bytes from the other program, and put them together. So, you might say that for the first two hundred and fifty we will use the mother's genes, for the second two hundred and fifty we will use the father's genes. Or, you might take randomly two hundred and fifty and set those to the father's values and two hundred and fifty others and set those to the mother's values. There are many different ways to do this and we are still exploring better techniques for this so-called sexual reproduction of computer programs. In the area of artificial life, we have had two international conferences now on artificial life.

A third way of getting in a new member of the colony is simply to create a totally random new member--we call that zapping, and this is also often done. So often you will be using cloning, sex, and zapping and you let the colonies evolve. Now, when I speak of colonies, it is an organizing technique in that we might have maybe a hundred creatures on the screen, but we might break those hundred creatures into twenty colonies, and each colony would have only five members and the cloning and sex would take place within the confines of each of the twenty colonies. Now, if you let this kind of program run for thousands of generations or hundreds of thousands of generations, typically we are looking at run times of at least several hours and preferably several days. On a fast computer, you will find really marked learning. You will find that the little creatures, if their task is to find each other and eat each others' trails, will be getting very good at it. If you have a maze that you want them to learn, they will get very good at finding their way though it. Essentially, they will evolve to satisfy any task you set them, if you can define it objectively so the fitness measure can take place automatically.

Now, this leads to an interesting issue. It is a case that is, I guess, somewhat analogous to biological patent issues, where if somebody invests maybe a year of computer time and evolves one of these colonies to a very high degree of fitness for performing some task, the output of this would be like the weights in a table of neural nets, or it would be the numerical values of the five hundred program bytes. So this would be something like a trade secret that your research expense had actually produced, and you would be able to use it. In itself, the secret would perhaps be something as simple as a page of numbers. It is simply that finding those numbers took a great deal of computer time.

What current tasks might artifical life programs be good for? Well, they have been used to help design circuit boards. In some sense, we could imagine animating electrons and thinking of the electrons as creatures; they try to crawl and get to the other components as rapidly as possible. We have done some things like that. Others have used a Connection Machine to find a maximally fast way for sorting a list of numbers--this has an artificial life representation.

The most promising area I am most interested in is designing systems to run robots. One of the on-going questions is why do we not have household robots yet, why do we not have robotic vacuum cleaners. The liability questions--the chances of the vacuum cleaner running over a baby--are so crippling that companies are very reluctant to put something out until it is extremely well tested. But it is partly an AI problem. The expert systems approach to having a robot trying to move around a room was carried out at Stanford--it would take the robot about nine hours to find its way across a crowded room. And then, they switched to an artificial life approach, and a robot can now find its way across the room very rapidly.

Generally speaking, the artificial life approach is appropriate when you have a programming situation where there are many independent parameters to set and there is little or no intuition about the best way to set them. In this case, we try a Darwinian approach--an evolutionary approach--and the essence of evolution is that we are going to have reproduction, sex and mutation, and natural selection. This tends to work pretty well. I believe in this approach because in my novel "Software" that's how I had the robots evolve to become free and intelligent. At this point, I am working to try to help make this reality.

In my remaining time, I want to talk briefly to you on the second topic that I mentioned in my abstract about artificial life and cyberspace. Normally when you do a computer graphic simulation, you limit it to the size of your screen which is 600 x 480 pixels or maybe 1,000 x 1,000 pixels. It is a really small arena that is two dimensional. What is needed to get richer computer simulations is to break through the computer screen and find yourself in a three dimensional world that is, in some sense, on the other side of the screen. The technology for that is just starting to come on line. There are really three new devices that are making this possible. One is the existence of what are known as "Z-buffered graphics boards", or as "pipelined graphics boards". The effect of a board like this is that you can download the geometry of a three dimensional scene, a scene such as this room with all the chairs in it, the ceiling lights and so on, and then you are able to use the mouse or your track ball to move your viewpoint around the scene as rapidly as a flying camera could move around. Usually there is a great deal of computation involved in changing a three dimensional viewpoint to figure out what is in front of what else so as to project a two dimensional image. But, with these new three dimensional graphics boards, we are able to effectively capture a three dimensional world in a bunch of chips.

There are two other elements to creating this illusion of a three dimensional computer world. Instead of looking at the screen you have the output go to two very small computer screens that you wear in a pair of goggles in front of your face. This gives you stereo vision, and there is a position sensor attached to your head. As you move your head, what you are seeing on the screens changes to match what you should see in the direction you are looking. So, you have the effect that you are inside this world looking around and you can look up behind you and see the complete surrounding of a three dimensional world filled up with three dimensinal computer graphics objects that might be moving around by programs that you have attached to them.

This second element gives a view into the three dimensional memory base that enhances the feeling that you are in it and the third element is to allow you the feeling that you are reaching into this three dimensional world. This is done with a variety of devices--the best known device is the Data Glove. The Data Glove is a nylon glove with optical fibres on it which sense how the bends in your finger join. You pull this glove on and as you move your hand, you project it into the three dimensional computer world. You see an image of your hand--this glove--floating there in a position just like your hand's position. This allows you to have a feeling of much more immersion in cyberspace.

As I said earlier, I am interested in helping to develop effective household robots. Rather than building actual mechanical prototypes of robots, we can imagine building cyberspace copies of robots. You would essentially put together a CAD model of a robot, of a machine that you wanted to have. You could put it into your virtual reality or cyberspace and you could watch how it behaves. This is a step beyond the first projected application for cyberspace which will be for architects to use as something that they can walk clients through. You can effectively be in the building before you have actually built it.

A final interesting possibility here would be if you had a lot of robots and were using artificial life techniques to evolve them to be better at getting around this virtual house. If a company wanted to stay a software company and not get involved in hardware, instead of selling the robots they could simply sell very detailed blueprints on how to put them together. They could sell virtual copies of it.

PROPRIETARY RIGHTS IN COMPUTER-GENERATED PRODUCTIONS

by

J. H. Reichman
Professor of Law
Vanderbilt University
Nashville, Tennessee
United States of America

I. ELECTRONIC INFORMATION AT THE OUTER EDGE OF WORLD INTELLECTUAL PROPERTY LAW

The specialized market for literary and artistic works operates with a peculiar legal and economic logic that, to varying degrees in different states, deliberately subordinates the price-setting function of the market to broader cultural policies. Copyright law thus occupies a privileged position among the laws that regulate trade in that its generous protective modalities need not strictly correlate with the demands of economic efficiency as measured in terms of utilitarian incentives to create.[1] At the same time, courts and legislators have traditionally ensured that the liberal treatment afforded artistic works would not undermine competition on the general products market, the operations of which are governed by the much stricter requirements of patent law. For this reason, the movement to bring industrial designs within the purview of the Berne Convention ultimately failed,[2] while borderline literary works of a functional nature--known as the "small change of copyright law"--normally obtain only thin protection against literal copying under the domestic laws of the industrialized countries.[3]

A. Systemic Anomalies of Applying Copyright Law to Computer-Generated Productions

The extension of copyright protection to computer programs from 1980 on, however, broke with this tradition,[4] and the resulting uncertainties[5] complicate the analysis of computer-generated productions. In this context, particular tensions arise when the copyright law's exclusive right to prepare derivative works is uncritically applied to digitized productions of every kind.[6]

1. The Derivative Work at Odds with Information Technologies

Under the United States Copyright Act of 1976, the exclusive right to prepare derivative works enables an author to recoup revenues generated from different uses of his or her work on each of the market segments where its expressive features are commercially exploited either in original or adapted form.[7] Strong protection of the author's market interest serves to align the "copyright" countries and the "authors' rights" countries without necessarily sacrificing the utilitarian ethos to which the former subscribe.[8] When factual or functional works become the object of infringement, however, courts instinctively narrow the scope of protection lest copyright owners acquire rights in the underlying facts or ideas that comprise the bulk of these works.[9] The judicial treatment of computer programs should logically conform to these precedents.[10] In reality, some recent decisions may have protected ineligible functional features[11] by grafting an overly broad reading of the right to prepare derivative works onto the sibylline definition of computer programs added to Section 101 of the Copyright Act in 1980.[12]

This provision defines computer programs in terms of a "set of statements or instructions to be used directly or indirectly in a computer in order to bring about a certain result."[13] Most computer programs evolve through revision, adaptation and transformation into an array of applications that are functionally "derived" from the set of instructions embodied in the programmer's initial solution.[14] By persuading courts to overextend the exclusive right to prepare derivative works, copyright owners can assert spurious proprietary claims to any subsequent innovations that to some extent exploit the original sets of instructions, even though all the relevant instruction sets contain little or no personal expression and fulfill primarily functional objectives.[15]

When courts then combine this pernicious version of the derivative work right with other protectionist doctrines, especially the "look and feel" test of copyright infringement used in some jurisdictions,[16] manufacturers of computer programs may sometimes obtain patent-like protection on soft conditions for a very long period of time.[17] These restraints on trade cannot be squared with the traditional justifications for a copyright system, and if the same doctrinal mix is blindly applied to computer-generated productions, it will compound the resulting social disutilities.

2. Electronic Art Tools

In principle, the traditional subject matters of copyright law should not suffer diminished levels of protection merely because digitization renders the application of technical doctrines more uncertain and complex.[18] For example, music publishers appear to be coping with the technology of sampling better than one might expect because samplers normally have an artistic and economic interest in presenting a recognizable fragment of another musical work. The recognizable quality of the sampled matter enables copyright proprietors to track the use and determine a reasonable fee.[19] It also discourages users of musical works from trying to evade the legal obligations of copyright law by appealing to technological freedom.

In contrast, sampling of the underlying sound recordings poses harder legal issues for the music industry,[20] especially in countries such as the United States that protect sound recordings in copyright law rather than under a neighboring rights regime.[21] Isolated sounds sampled from particular sound recordings partake even less of personal creative expression than sound recordings as a class,[22] and the sampler may so process the sampled bits of sound that they cannot be attributed to any identifiable source.[23] In this state, the isolated sounds taken from the original recording function mainly as a sound specimen that merges with other specimens into units or structures of sound corresponding to the form that the processor imposes upon them. Arguably, even these sound specimens possess commercial value or the samplers would not trouble to appropriate them.[24] Yet, one balks at the prospects of drawing meaningful distinctions between the universe of free sounds in which any sound engineer can freely operate and a putative universe of protected sound from which an engineer is excluded merely because the sounds in question happen to emanate from identifiable human beings.

The United States copyright law, though embracing sound recordings, limits their scope of protection to actual dubbing and permits unauthorized simulation or imitation.[25] Nevertheless, by treating sound recordings as original works of authorship within Section 102(a), the 1976 Act invests the "author"--normally a corporation--with the exclusive right to prepare derivative works.[26] Paradoxically, the immunities set out in Section 114(b) further weaken the sampler's legal position because the very language that allows third parties freely to simulate the sounds on a phonorecord also prohibits the making of an unauthorized derivative work that rearranges, remixes or otherwise alters the actual sounds of a protected sound recording.[27] Although this language may produce unintended effects in the realm of digital sampling, it nonetheless encourages copyright proprietors to lodge harassing claims even when the sampled sounds are processed beyond recognition.

The processing of sampled sounds into their final form usually requires a high degree of technical skill. Use of the latest equipment may indeed require as much skill and even "authorship-like" qualities as can be attributed to the sampled sound recording itself. While this suggests that users may sometimes qualify as authors or joint authors of the end product,[28] characterization of the digitizing equipment as an art "tool" tends to foreclose certain proprietary claims that manufacturers of the equipment might otherwise be tempted to launch.

The distinction between a work of art and an art tool is usefully illustrated in regard to digitized reproductions of artistic works. Art reproductions have always posed a formal paradox because the indispensable quantum of originality that copyright law imposes seems negated by the need to make a perfect copy of a pre-existing work. Courts have tried to resolve this paradox by stressing the dependence of good reproductions on high artistic skills and by insisting that the reproducer cannot altogether mask the imprint of his or her own personality, at least not from other skilled experts in the trade.[29] However, the introduction of digitized equipment to perfect the reproduction weakens the case for copyright protection by further lessening the quantum of artistic skill needed to break out of the preexisting paradox.

At the same time, digitized reproduction equipment depends on computer programs, and the open-ended definition of computer programs in Section 101[30] lacks any intrinsic limits like those imposed by the form of a novel or a song. Courts have already used this definition to justify derivative rights in object code programs that were mechanically "translated" from source code or from earlier embodiments of the programmer's solution.[31] Unless courts more aggressively test the definition of computer programs against the criteria of eligibility set out in Sections 102(a) and 102(b),[32] this reasoning permits digitized art reproduction to be characterized as derivative manifestations of the set of instructions embodied in the code,[33] an argument that could vest authorship in the programmer or his employer.[34]

Such claims remain counterintuitive in the domain of graphic art and would seldom succeed in the end. Although digital processing greatly reduces or eliminates the margin of error and augments the user's ability to modify and manipulate the underlying works,[35] most electronic art tools still require the skills of an artist to perfect the end product. This quantum of artistic creation should continue to attract copyright protection in its own right whenever it can be distinguished from the contribution of the electronic art tool for purposes of eligibility.

Nevertheless, the quality of a digitized art reproduction or of a colorized black and white film--and thus its desirability in the market--is significantly determined by a computer program, and the instruction sets embodied in the program are to some extent mirrored in the reproduction along with the user's own skilled contribution.[36] Although the copyright law has never permitted manufacturers of artistic tools, such as paints and paint brushes, to advance derivative claims to the works resulting from their use, pressures for different results could grow as the clarity of the distinction breaks down and the tools acquire, so to speak, a mind of their own that contributes significantly to the end result.

Professor Davis inevitably poses this dilemma in his discussion of Aaron's machine-produced paintings, which are largely a product of aesthetic principles "taught" to the machine by its originator.[37] Similar developments reported by Professor Samuelson include a system that generates "synthesized music in response to the movements of a person (or other forms) located at the vortex of two video cameras" and another system that converts visual images into sounds that enhance the viewer's ability to interpret the visual image under examination.[38] In one application of the latter system, a cursor reads a photograph of human tissues and emits sounds that differentiate healthy from cancerous images.[39]

For the moment, Aaron's art machine constitutes a limited case in that it does posses the capacity to produce new works of art with few or no inputs from the user. One can therefore envision a nexus between the programmer's original expression as embodied in the computer program and his or her status as author or co-author of the resulting artistic work.[40] Ownership problems would then be relegated to specific licensing agreements between the owner of the machine and would-be users.

In other situations, however, manufactured legal difficulties can be lessened by holding to the concept of the tool and by asking for what purpose the tool is being used. In the first of Samuelson's examples as set out above, the digital equipment enables users to produce a form of music (or a cluster of sounds) by changing the movements of the body or limbs. Although the totality of sounds is, perhaps, encompassed in the computerized instructions buried in the equipment, the purpose of the equipment is to enable purchasers, such as handicapped people, to turn motion into a certain kind of music. None of the cultural or economic policies of the copyright law are served by extending the proprietary claims of the toolmaker to the artistic product in this case, and some policy goals could be disserved if copyright law too readily begins to treat the makers of art tools as authors of derivative works.[41]

The case of the tissue-reading sound equipment seems even farther removed from the proper scope of artistic copyright as such. There is no personal intellectual creation in the resulting sounds; there is simply a more or less accurate diagnostic test that results from using a medical tool. In such cases, joint and derivative work claims raise false issues because the market determined value of the output has nothing to do with the personal expression of the creator as commercially exploited on different market segments.[42] The real issue is that medical systems of this kind, or their most commercially valuable components, may not qualify for patent protection,[43] and the tool manufacturer could therefore lack sufficient lead time in which to recoup his investment or to warrant undertaking the innovation in the first place.[44]

In this respect, applying the tool concept to the problems of digitization suggests that an exclusive right to control end use--drawn from industrial property law--might help the electronic toolmaker far more than some spurious right to prepare adaptations of artistic works.[45]

To the extent that most electronic art tools continue to remanage and rearrange existing artistic works, as in the case of music sampling machines and art reproduction machines, the derivative work rights of the underlying authors deserve protection because the computer-generated outputs actually embody and exploit their original expression on a different market segment.[46] The tool maker, in contrast, will normally have facilitated such a transformation without adding personal expression of his own. In that event, he should no more be entitled to a proprietary share in the user's final artistic product than the makers of paint brushes and canvas in less technologically refined epochs.[47]

The importance of this line of analysis becomes even clearer the moment one turns from present-day problems raised by electronic art tools to the market distortions that ensue when more utilitarian digital processing tools are admitted to full copyright protection on the excuse that they "convey information."[48]

3. Electronic Information Tools

Comparative intellectual property law suggests other legal models that are better suited economically and conceptually to addressing the problems posed by digitalized information processing than the derivative rights doctrine of copyright law. Of particular interest in this regard is the protection that utility model laws historically conferred upon the external configurations of certain handtools and other everyday implements.[49]

Functional improvements in early handtool design, especially that of agricultural implements, almost always entailed strong elements of form or shape that were rarely inventive enough to qualify for patent protection.[50] Like ornamental designs, these handtool designs were embodied in products sold on the open market, which made them ineligible for trade secret protection. Legislative decisions to protect utility models thus recognized that handtool designs remained as vulnerable to appropriation by third parties as ornamental designs of useful articles that were protected under sui generis design laws.[51] Both design categories evolve through incremental innovation, and the physical support bears the designer's know-how on its face in either case.[52] Because most ornamental designs laws exclude functionally determined designs by definition,[53] utility model laws aimed to plug a gap in the intellectual property universe that has widened inordinately with the advent of important new technologies.[54]

Utility model laws operate with a stricter legal discipline than that of the sui generis design laws, a phenomenon usually ascribed to the functionality of the designs they protect.[55] For example, utility model laws require a qualitatively significant innovation in exchange for a short term of immunity from competition.[56] While utility model laws nominally confer a bundle of exclusive rights comparable to that of patent law, they exempt functional designs from a full examination of the prior art and provide a narrow scope of protection in keeping with "the limited character of the invention."[57]

Because utility model laws are expected to protect functional improvements attained by means of specific shapes,[58] they cannot protect processes, and until recently, they did not, for this reason, apply to most electronic circuit designs.[59]

Utility model laws follow the principle of exhaustion, in common with most developed patent systems.[60] This meant that, once a protected handtool was sold on the open market, manufacturers normally retained no further control over the uses to which their innovative functional designs were put.[61] Utility model laws thus required innovative toolmakers to exact the reward for their products in the monopoly prices applicable at the time of first sale; but they normally precluded manufacturers from asserting any claim to the value added to other products by those who purchased and used their tools. In this way, the legal protection of handtool designs implicitly recognized that users of the tools add significant value to their own products, an enterprise in which the toolmaker contributes little beyond the efficiencies that entitled him or her to protection in the first instance.

Commentators often criticize utility model laws for unjustifiably undermining the integrity of the patent laws.[62] Their true shortcoming is that they arbitrarily enabled only certain categories of industrial designers to protect functional product configurations while ignoring a general need to protect unpatentable, noncopyrightable embodiments of know-how that could not qualify for trade secret protection.[63] On the one hand, utility model laws degenerate into petty patent laws that now sometimes protect electronic circuit designs.[64] On the other hand, utility model laws were precursors of a bevy of hybrid legal solutions that are currently being thrown at new technologies with increasingly unsatisfactory results.[65]

This writer's general thesis, expounded in other fora, is that the world's intellectual property system needs a single law to protect applied scientific know-how as such.[66] On the whole, a modified patent model has not solved the problems of protecting industrial designs generally, and it seems doubtful that it would adequately fill the needs of a new legal paradigm.[67] Rather, there are good reasons why modern industrial know-how tends to seek the copyright paradigm,[68] and the most promising approach to the protection of new technologies as a whole appears to reside in a modified copyright approach.[69]

A proper know-how law along these lines could be shaped to resolve many of the proprietary complications that computer-generated productions are likely to elicit. Absent such a law, the intellectual property community should look to the tool model of industrial property law in approaching computer-generated productions and not to the "derivative work" concept of artistic property law.[70]

The task is to ensure that the creators of electronic information tools receive adequate incentives without requiring users to recognize derivative ownership rights in posterior innovation made with their tools and without impairing the ability of second comers to enhance the efficiency of these same tools. At times, a lump sum sales price will adequately compensate the toolmaker for his efforts. At other times, the nature of the electronic information tool may require per use payments to determine its true market value. None of these remuneration problems become insuperable so long as the goal remains that of compensating the tool maker for the behavioral impact of the tool itself and not for putative "adaptations" of its constituent parts that result from using the tool for its intended purpose.

The tool concept also sheds light on the uncertain status of small-change literary works in the pre-digital world of copyright law,[71] and it helps to solve problems raised by the advent of computer-aided design. The small change always strained general principles of copyright protection because they were essentially low authorship, functional tools hiding out in artistic property law.[72] In this sense, copyright protection of computer programs as literary works merely converted the "small-change" anomalies of yesterday into the "big bucks" anomalies of today, anomalies that result when courts and legislators are pressured to tolerate parochial derogations from the competitive ethos. Properly conceptualized, copyright law should never afford functional works more than "thin" protection against wholesale appropriation of surface expression.[73] And there is no market interest to trigger derivative work rights so long as users apply functional works to the purposes for which they were intended[74] or second comers take only utilitarian features without duplicating surface expression.[75]

As applied to computer-aided design, these same principles limit the toolmakers' capacity to claim derivative rights in their users' end products. No matter how refined the computer-aided technology may be, a user's output usually competes on the products market and not on the market for artistic works.[76] What counts on the tool segment of the products market, moreover, is not normally the personal expression of either the toolmaker or the user, but rather the potential value to be added by the user who applies the tool to the task of making better products at more competitive prices.[77] In principle, the toolmaker's reward should come from his sale price or licensing fees, and not from the value-adding uses to which his innovative tool is put.[78]

It follows that the outputs of computer-aided design, once allocated to the user of the tool and not to its maker, must seek protection according to their variable natures under whatever intellectual property regimes happen to apply. The computer-aided design of a sweater, for example, might qualify for protection under the ornamental design laws of the Benelux countries; under the new, unregistered design right of the United Kingdom; under copyright law in France; or even under increasingly protectionist applications of Section 43(a) of the United States Trademark Act to unregistered "appearance trade dress."[79] However, courts should resist proprietary theories seeking to link the outputs of computer-aided design with the inputs of a manufacturer's computerized tool so long as the programmers of the tool add comparatively little value to the user's end product in relation to that added by the user himself.

B. Artificial Intelligence as Semi-Autonomous Know-How Machines

The analytical framework outlined above yields less certain results in the domain of artificial intelligence systems because these systems are, by definition, capable of semi-autonomous forms of discernment.[80] One can already envision situations in which both the makers and the users of such machines will keep adding inputs in order to achieve system-determined outputs that greatly exceed the individual contributions of either party. At the limit, while both maker and users continue to input valuable information, the artificial intelligence machine may itself add the bulk of any new or additional value to previously available outputs.[81]

Investigation of these phenomena is greatly handicapped by the lack of firm and workable decisions about the patenting of computer programs and of program-related inventions, and by the skewed empirical results that have ensued from this lack of consensus.[82] To the extent that computer programs become so firmly anchored in the patent law that algorithms obtain direct or indirect protection, this body of law will probably undergo significant adjustments to prevent undue blockage of scientific progress.[83] Meanwhile, a law to protect applied scientific know-how remains beyond the horizon, and there is little experience to suggest where its outer limits might lie.

In principle, an adjusted patents framework could be applied to artificial intelligence machines, even if the prospect of monopolizing self-executing algorithms capable of discerning choices would greatly compound the public policy dilemmas already troubling computer science.[84] The more that patents on artificial intelligence machines tended to collapse the distinction between process and utility patents, for example, the more that patentees would try to include outputs of these machines within the scope of their claimable inventions. Broadening the doctrine of patent misuse might then become necessary to limit monopolistic control over activities on which both scientific and technological progress could depend.[85]

Over time, even the most astounding breakthroughs in artificial intelligence may give way to more routine applications of basic principles, as occurred in biotechnology once the technical dimensions of the recombinant DNA breakthrough became better understood.[86] For example, progress in artificial intelligence requires computer science to standardize programs for use as building blocks in larger, more powerful systems. These components would presumably shed their novelty along the way. As important segments of commercially valuable systems--including outputs--failed to qualify for patent protection despite subject-matter eligibility, new and highly disruptive efforts will be made to extend copyright law[87] to every technological innovation that bears its know-how on its face.[88]

Meanwhile, copyright law protects the static components of a computer program for an exorbitant period of time without directly reaching the dynamic configuration that accounts for the commercially valuable behavior of any given system. If patent law gives too much protection to too few program-related inventions, present-day copyright law provides far too much protection to far too many program features. These tendencies hinder progress in the art without adequately rewarding the applied scientific know-how on which it depends.[89]

To the extent that artificial intelligence consists of computer software, as Professor Davis insists,[90] stuffing artificial intelligence into copyright law merely exacerbates all the unsolved problems afflicting the legal protection of software in general. Questions of authorship and ownership, for example, are complicated enough when patent, copyright, and trade secret laws apply concurrently to essentially the same software innovation, especially if it results from the collective efforts of a team of investigators working on the same problem.[91] Copyright protection of artificial intelligence will only magnify these difficulties and render the lines separating creator from producers and users more uncertain.[92] Too many potential owners arguing over too many potentially valuable interests will then further retard the pace of innovation independently of other anti-competitive effects inherent in the application of artistic property law to functional designs.[93]

Before these and other related problems strain world intellectual property law beyond the breaking point, it would be helpful to have a proper know-how law on the books and even better if its legal machinery were fully elaborated and its limitations well understood. A proper know-how law would aim to protect the functional behavior achieved by means of certain aggregates of information and not the aggregates of information themselves. Such a law might be capable of protecting functional improvements obtained by rearrangements of pre-existing components. And it could conceivably protect outputs of artificial intelligence machines without necessarily succumbing to pernicious extensions of either patent or copyright doctrines that were devised for an era in which the distinction between theoretical and applied science still made perfect sense.[94]

C. Beyond Know-How: Outputs of Advanced Artificial Intelligence Systems

Once the tool model that worked fairly well in the context of computer-aided design begins to break down, however, claims linking producers and users in terms of values generated by their interactive processing of information could support an hypothesis that electronic information--digital data structures--embodied in the system and its outputs should become an object of legal protection in its own right.[95] In the realm of artificial intelligence, in other words, effective legal protection could eventually depend upon a willingness to protect aggregates of information as such and not just the functional configuration representing the know-how responsible for particular systemic behavior.[96] If this hypothesis proves valid, it would confirm insights about a linkage between current problems of protecting computerized data bases and future problems of protecting the outputs of artificial intelligence systems.[97]

The possibility that electronic information might one day constitute a marginal case beyond even the marginal case of applied scientific know-how is not an hypothesis this writer delights in airing. There is already too much loose talk about "information" that adds little to legal or economic analysis. Such talk also deflects attention away from a proper know-how law that could integrate the proliferating array of legal hybrids into a unified field of protection.[98] Until the intellectual property community attains this goal and digests the ensuing results, one prefers not to contemplate additional protective schemes that could burden free research or encourage oligopolistic industries to foster new barriers to entry.

Nevertheless, within the confines of a worldwide symposium that is deliberately and collectively peering into its crystal ball, scientists and legal scholars must face up to the remote but scary possibility that electronic information, and not just know-how, could one day become a fit subject of intellectual property law in a world populated by artificial intelligence machines. If and when artificial intelligence begins to provide a stream of semi-autonomous decision-making tools, in short, electronic information may have to be viewed as a kind of tool in its own right. In that event, those primitive, nineteenth century utility model laws may still have other valuable lessons to impart.

II. NEED FOR A NEW LEGAL PARADIGM

This author's previous studies have tried to show the extent to which the world's intellectual property system has been overwhelmed by new technologies that do not fit within the patent and copyright systems. Problems arise mainly because investors in incremental innovation bearing know-how on its face lack artificial lead time in which to recoup their investment and turn a profit. The solution requires a law to protect applied scientific know-how as such, regardless of the medium of expression, that deters free riders without unduly impeding fair followers from developing incremental innovation of their own. The evidence suggests that such a law should be built on modified copyright principles without, however, succumbing to the philosophical mystique of the mature copyright paradigm, whose powerful reproduction rights and long term of protection implement cultural policies that are largely irrelevant to the needs of a competitive market.

The present study, like its forerunners, tends to confirm the disarray that results from concurrent application of several intellectual property laws to new technologies for which they were not devised. While some may continue to believe that more of every kind of protection benefits innovators, the evidence suggests that the legal process itself has begun to slow the pace of innovation and that it has especially harmed the small and medium-sized firms that are--or were--its major exponents.

Wherever one looks, the tendency is to disaggregate each new technology into its component parts and then to assimilate these parts to existing legal paradigms, with the excuse that international politics justifies the ensuing distortions to domestic market forces. The end result is to enmesh each of these often trivial components in endless ownership and scope of protection issues that each of the competing legal subcultures logically spawns. Meanwhile, the specific functional and behavioral impact achieved by innovators working in the new technologies is nowhere adequately protected as such. And no specific set of authorship and ownership rules can be devised to implement a coherent set of policy goals because the policy goals actually being implemented pertain to semi-obsolete or largely irrelevant forms of legal protection.

The more that these overlapping legal subcultures compete with each other, the harder it becomes to disentangle applied scientific know-how from their separate protective strands, and the more the world's entire intellectual property system lurches towards incoherence. Unless the urge to throw existing legal regimes at a moving target is resisted in the interest of a more rational and constructive debate, the advent of still more difficult challenges in the form of computer-generated works, computer-aided design, and artificial intelligence machines will bring the system to a halt.

Some may fear that a sui generis exploration of applied scientific know-how would lead to even worse complications and greater uncertainties in the future. They should look around them to see what is actually happening to existing intellectual property regimes that are increasingly put to abusive and inconsistent applications. Wherever one looks, indeed, one is struck by the extent to which the laws applicable to patents, copyrights, trade secrets, unfair competition, trade marks, and industrial design are increasingly destablized by the need to deal with aspects of new technologies for which they are inherently unsuited. Rather than facing up to the new problems that might arise from a single, *sui generis* regime to protect applied scientific know-how, in short, the intellectual property community is currently experiencing the simultaneous evolution of six or more poorly designed *sui generis* laws, as each traditional regime mutates in unexpected ways under the pressure of events.

In the long run, unless a law to fill the gap is placed on the drawing board in a timely fashion, the risk is that the entire system will discredit itself in the eyes of neutral economic policymakers. In that event, those who have historically opposed the rights of authors and inventors in the name of totally free trade may yet score a victory that was unthinkable at the height of the nineteenth century patents controversy.

NOTES

1 See, e.g., Reichman, Goldstein on Copyright Law: A Realist's Approach to a Technological Age, 43 STAN. L. REV. 943, 946-49 (1991) [hereinafter cited as Realist's Approach]. For example, incentive theory cannot account for the moral rights. Nor will it adequately explain such paternalistic measures in American copyright law as the right to terminate transfers, see 17 U.S.C. Secs. 203, 304(c) (1988), or even the long period of protection, which enables living authors and their immediate heirs to partake of revenues generated many years after the creation of their works. See, e.g., Limperg, Duration of Copyright Protection, 103 R.I.D.A. 53, 68-69, 72-77 (1980). For the dominant role of the incentive rationale in United States copyright law, see generally 1 P. GOLDSTEIN, COPYRIGHT: PRINCIPLES, LAW AND PRACTICE 3-54 (1989).

2 See generally Reichman, Design Protection and the New Technologies: The United States Experience in a Transnational Perspective, 19 BALTIMORE L. REV. 6, 126-36 (1990), abridged version reprinted in 1991 INDUSTRIAL PROPERTY (May (Pt. 1), June (Pt. II)) [hereinafter cited as Designs and New Technologies].

3 See, e.g., ADOLF DIETZ, COPYRIGHT LAW IN THE EUROPEAN COMMUNITY 30-33 (1978). For the "thin" copyright doctrine in the United States, see Feist Publications v. Rural Telephone Service Co., 111 S.Ct. 1282, 1294-95 (1991); Reichman, Realist's Approach, supra note 1, at 970-73.

4 See, e.g., Samuelson, CONTU Revisited: The Case Against Copyright Protection for Computer Programs in Machine-Readable Forms, 1984 DUKE L.J. 663; Goldstein, Infringement of Copyright in Computer Programs, 47 U.PITT.L.REV. 1119 (1986).

5 See generally, COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD, NATIONAL RESEARCH COUNCIL, NATIONAL ACADEMY OF SCIENCES, INTELLECTUAL PROPERTY ISSUES IN SOFTWARE 43-93 (1991); Last Frontier Conference Report on Copyright Protection of Computer Software, 30 JURIMETRICS J. 15 (1989) (consensus statement by ten law professors) [hereinafter cited as Conference Report].

6 See, e.g., Samuelson, Digital Media and the Changing Face of Intellectual Property Law, 16 RUTGERS COMPUTER & TECH. L.J. 323, 328-34 (1990) [hereinafter cited as Digital Media].

7 See 17 U.S.C. Secs. 101, 103, 106(2) (1988); 2 P. GOLDSTEIN, supra note 1, at 32-37.

8 See, e.g., S. RICKETSON, THE BERNE CONVENTION FOR THE PROTECTION OF LITERARY AND ARTISTIC WORKS: 1886-1986, at 286-87, 293-95, 398-400 (1987); WORLD INTELLECTUAL PROPERTY ORGANIZATION, GUIDE TO THE BERNE CONVENTION FOR THE PROTECTION OF LITERARY AND ARTISTIC WORKS (PARIS ACT 1971) 76-77 (C. Masouyé & W. Wallace eds. 1978); 2 P. GOLDSTEIN, supra note 1, Secs. 7.1.2, 7.3.2.

9 See, e.g., Feist Publications v. Rural Telephone Service Co., 111 S.Ct. 1282 (1991); Baker v. Selden, 101 U.S. 99 (1879).

[10] See, e.g., Reichman, Computer Programs as Applied Scientific Know-How: Implications of Copyright Protection for Commercialized University Research, 42 VANDERBILT L. REV. 639, 684-87, 689-97 (1989) [hereinafter cited as Programs as Know-How]; Nimmer, Bernacchi & Frischling, A Structured Approach to Analyzing the Substantial Similarity of Computer Software in Copyright Infringement Cases, 20 ARIZ. ST. L.J. 625, 628-34, 642-49 (1988), adapted and reprinted in 3 M. & D. NIMMER, NIMMER ON COPYRIGHT / 13.03 (rev. ed. 1990); Menell, An Analysis of the Scope of Copyright Protection for Application Programs, 41 STAN.L.REV. 1045 (1989).

[11] See, e.g., Whelan Assocs. v. Jaslow Dental Labs, 797 F.2d 1222 (3d Cir. 1986) (broad copyright protection for elements of structure, sequence, and organization), cert. denied, 479 U.S. 1031 (1987); Apple Computer, Inc. v. Franklin Computer Corp., 714 F.2d 1240 (3d Cir. 1983); Lotus Dev. Corp. v. Paperback Software Int'l, 740 F.Supp. 37 (D. Mass. 1990) (protecting menu command structure of computer spreadsheet program, including the choice of command terms, the structure and order of these terms, and long prompts, despite evidence of standardization and functional efficiencies); Manufacturing Technologies Inc. v. Cams, Inc., 706 F.Supp. 984 (D. Conn. 1989); Digital Communications Assocs. v. Softklone Distributing Corp., 659 F.Supp. 449 (N.D. Ga. 1987). But see Plains Cotton Cooperative v. Goodpasture Computer Service, 807 F.2d 1256 (5th Cir. 1987), cert. denied, 108 S.Ct. 80 (1987) (sequence and organization not protectible when largely determined by market factors).

[12] See Software Protection Act of 1980, Pub. L. No. 96-517, Sec. 10(b), 94 Stat. 3015, 3028 (codified as amended at 17 U.S.C. Secs. 101, 117 (1988)).

[13] 17 U.S.C. Sec. 101 (1988).

[14] See, e.g., Spoor, Expert Systems and Copyright, in G.P.V. VANDENBERGHE et al., ADVANCED TOPICS OF LAW AND INFORMATION TECHNOLOGY 93, 101-03 (1989); Gemignani, Copyright Protection: Computer-Related Dependent Works 15 RUTGERS COMPUTER & TECH. L.J. 383, 383-87 (1989).

[15] See, e.g., Samuelson, Allocating Ownership Rights in Computer-Generated Works, 47 U. PITTSBURGH L. REV. 1185 (1986) [hereinafter Allocating Rights]; Goldstein, General Report: Computer-Assisted and Computer-Generated Creation of Literary and Artistic Works, in L'INFORMATIQUE ET LE DROIT D'AUTEUR 439, 442-43 (1990) (proceedings of the ALAI Congress, Quebec, Canada, Sept. 26-30, 1989) [hereinafter General Report].

[16] For the different tests of copyright infringement that U.S. courts currently apply, see Reichman, Realist's Approach, supra note 1, at 957-58, 957 nn. 94 & 95. For a thoroughgoing and accurate explanation of these tests and how they should be applied, see generally 2 P. GOLDSTEIN, supra note 1, at 3-38.

[17] See, e.g., Menell, supra note 10, at 1082 (concluding that "the Whelan rule makes it difficult for others wishing to market programs performing the same task as the first comer to perform it as effectively" and thus "enables first comers to 'lock up' basic programming techniques"). See also Karjala, Copyright, Computer Software, and the New Protectionism, 28 JURIMETRICS J. 33, 34-36, 62-96 (1987).

[18] See generally, Goldstein, General Report, supra note 15, at 441-48.

[19] See, e.g., Biederman, "Licensing of Sampled Musical Works and Sound Recordings" paper presented to the Mid-Year Meeting of the Copyright Society of the U.S.A., Nashville, TN, Feb. 7-9, 1991.

[20] See, e.g., McGraw, Sound Sampling Protection and Infringement in Today's Music Industry, 4 HIGH TECH. L.J. 147 (1989); Note, Digital Sound Sampling, Copyright and Publicity: Protecting Against the Electronic Appropriation of Sounds, 87 COLUMB. L. REV. 1723 (1987).

[21] See 17 U.S.C. Sec. 101 (1988) (defining sound recordings as "works"); id. Sec. 102(a)(7) (listing sound recordings as works of authorship). But see 17 U.S.C. Sec. 106(4) (1988) (not conferring exclusive public performance right on sound recordings); id. Sec. 114(a) (confirming denial of public performance right); Id. Sec. 114(b) (limiting reproduction right to duplication but not imitation or simulation). The Berne Union countries have traditionally stigmatized the level of authorship in sound recordings as insufficient to warrant copyright protection. Neighboring rights laws, in turn, characteristically protect sound recordings only against duplication and not against imitation or simulation. See generally 2 P. GOLDSTEIN, supra note 1, at 686-88, 701-04; WORLD INTELLECTUAL PROPERTY ORGANIZATION, GUIDE TO THE ROME CONVENTION AND TO THE PHONOGRAMS CONVENTIONS (1981).

[22] See 17 U.S.C. Sec. 102(a) (1988) (requiring original works of authorship); Feist Publications v. Rural Telephone Service Co., 111 S.Ct. 1282, 1289 (1991) (requiring minimum degree of creativity to meet constitutional mandate). For the growing convergence between domestic and foreign law in regard to the originality requirement, see Reichman, Realist's Approach, supra note 1, at 951-55 (citing authorities).

[23] See, e.g., McGraw, supra note 20.

[24] See, e.g., Goldstein, General Report, supra note 15, at 449-50 (comparing the sampling of sound fragments to library photocopying of entire works in terms of common difficulties of policing infringement, but fearing that "in the aggregate, both forms of appropriation may seriously erode economic and moral rights in original works").

[25] See supra note 22. In this respect, domestic law emulates the neighboring rights laws of the European Community countries. See supra note 21 and accompanying text.

[26] See 17 U.S.C. Secs. 101, 102(a)(7), 103, 106(2), 201 (1988); Note, supra note 20, at 1744; supra note 22. A priori, record companies seeking to enforce their adaptation rights should expect to encounter several defenses likely to benefit those who sample isolated sounds from a copyrighted sound recording. For example, the lifting of a small component of sound or of separate but small sound bites might amount to an insubstantial or de minimis appropriation and thus fail one of the subtests of the substantial similarity doctrine. If the sampled sounds as ultimately reprocessed bear little or no recognizable similarity to the sounds as used on the original recording, then the sampler's productions cannot be deemed to displace the market for the original sound recording. The defendant's reprocessed sounds could thus fail the "audience" subtest of the infringement calculus as some leading courts

apply it. See, e.g., 2 P. GOLDSTEIN, supra note 1, at 7. In this way, the incentive theory of United States copyright law pulls against protection of sampled sound recordings in certain instances where natural rights theory may tilt the other way.

The availability of unfair competition law to redress the balance remains uncertain and controversial. See 17 U.S.C. Sec. 301 (1988). Nevertheless, to the extent that certain sounds function as an unregistered aural service mark, there may be some valid basis for avoiding a likelihood of confusion under the federal law of unfair competition. See Lanham Act, ch. 540, Sec. 43(a), 60 Stat. 427, 441 (1946) (codified at 15 U.S.C. Sec. 1125 (1988)) as amended by Sec. 132 of the Trademark Act of 1988. Cf. 17 U.S.C. Sec. 301(d) (other federal laws not preempted).

27 See 17 U.S.C. Sec. 114(b) (1988); supra note 22.

28 See, e.g., Goldstein, General Report, supra note 15, at 445-47; cf. Spoor, supra note 14, at 101-03.

29 See, e.g., Alfred Bell & Co. v. Catalda Fine Arts, Inc., 191 F.2d 99 (2d Cir. 1951); Alva Studios, Inc. v. Winninger, 177 F.Supp. 265 (S.D.N.Y. 1959).

30 See supra note 13 and accompanying text.

31 See, e.g., Apple Computer, Inc. v. Franklin Computer Corp., 714 F.2d 1290 (3d Cir. 1983); see also Apple Computer, Inc. v. Formula Computer Corp., 725 F.2d 521 (9th Cir. 1984).

32 See 17 U.S.C. Secs. 102(a) (1988) (originality, creativity, and work of authorship), 102(b) (excluding ideas, methods, systems, principles, processes and discoveries).

33 See, e.g., Goldstein, General Report, supra note 15, at 444-46 (discussing German and United States precedents).

34 17 U.S.C. Sec. 201(a) & (b) (1988). See generally Reichman, Programs as Know-How, supra note 10, at 673-81. Such arguments are buttressed by joining the open-ended definition of computer programs with the increasingly elastic treatment of "derivative works" and "joint works" in result-oriented legal opinions. See most recently Reichman, "Proprietary Rights in the New Landscape of Intellectual Property Law: An Anglo-American Perspective," paper delivered to the ALAI's Congress of the Aegean Sea II, Athens, Greece, April 19-26, 1991 [hereinafter cited as Proprietary Rights].

35 See, e.g., Samuelson, Digital Media, supra note 6, at 328-33.

36 While the European Community countries appear to lean towards joint authorship in this situation, United States law arguably tends to view "the user's work ... as a derivative work, of which the user will be the sole author if he used the program lawfully." Goldstein, General Report, supra note 15, at 447.

[37] See Davis, "The Assumptions Are Broken," paper delivered to WIPO's Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence, Stanford University, Palo Alto, CA, March 25-28, 1991.

[38] Samuelson, "Some New Kinds of Authorship Made Possible by Computers and Some Intellectual Property Questions They Raise," paper presented to the Intellectual Property and Authorship Conference, Case Western University, Cleveland, Ohio, April 19-21, 1991).

[39] Id.

[40] See, e.g., Samuelson, Allocating Rights, supra note 15, at 1200-05.

[41] For the view that the overriding policy goal of copyright law is "to encourage the widest possible production and dissemination of literary, musical and artistic works," see 1 P. GOLDSTEIN, supra note 1, at 11. For a discussion of other views, see Reichman, Realist's Approach, supra note 1, at 947-49, 978-79.

[42] See supra text and authorities accompanying notes 7-11.

[43] Ineligibility may arise either because the computer-related innovation is viewed as nonpatentable subject matter or because it fails to meet the threshold requirements, especially that of nonobviousness under 35 U.S.C. Sec. 103 (1988). For contrasting views in regard to patentability, compare Samuelson, Benson Revisited: The Case Against Patent Protection for Algorithms, and Other Computer-Program-Related Inventions, 39 EMORY L.J. 1025 (1990) with Chisum, The Patentability of Algorithms, 47 U.PITT.L.REV. 959 (1986). For recent developments, see Stern, Tales from the Algorithm War: Benson to Iwahashi, It's Déjà Vu All Over Again, 18 AIPLA Q. J. 371 (1991); Laurie, "The Patentability of Artificial Intelligence Under U.S. Law," paper presented to the WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence, Stanford University, Palo Alto, Calif., March 25-27, 1991.

[44] See generally Reichman, Programs as Know-How, supra note 10, at 652-662.

[45] See 35 U.S. Sec. 271 (1988) (conferring exclusive right to make, use or sell the patented invention). The exclusive right to control end use is not granted copyright owners by 17 U.S.C. Sec. 106 (1988); see, e.g., Brown, Eligibility for Copyright Protection: A Search for Principled Standards, 70 MINN. L. REV. 579, 588-89 (1985). Nevertheless, the status of such a right remains controversial in both domestic and foreign law. See id., at 590-98; Reichman, Programs as Know-How, supra note 10, at 693, 693 n. 288 (citing authorities).

[46] See, e.g., Seecof, Scanning Into the Future of Copyrightable Images: Computer-Based Image Processing Poses a Present Threat, 5 HIGH TECH. L. J. 371 (1990) (discussing digital processing of photographic works); Biederman, supra note 19.

[47] See, e.g., Goldstein, General Report, supra note 15, at 445-46. But see id. at 446-47 (recognizing other situations in which both programmer and user contribute substantial creative elements to true literary or artistic products, and noting that in such cases, both derivative work theory and joint work theory may properly be used, with results that vary from one country to another).

48 See 17 U.S.C. Sec. 101 (defining "useful articles" to exclude an article "having an intrinsic utilitarian function that is not merely to portray the appearance of the article or to convey information") (emphasis supplied). This definition facilitates the "unity of literature" approach to computer programs adopted in the very same section of the 1976 Act that rejects the "unity of art" approach to industrial designs. See 17 U.S.C. Sec. 101 (1988) (definition of "pictorial, graphic, and sculptural works"). See generally Reichman, Designs and New Technologies, supra note 2, at 8-12.

49 See, e.g., Higashima & Ushiku, A New Means of International Protection of Computer Programs Through the Paris Convention--A New Concept of Utility Model, 7 COMPUTER L.J. 1, 15-22 (1986); Note, Petty Patents in the Federal Republic of Germany: A Solution to the Problem of Computer Software Protection?, 8 SW. U. L. REV. 888-909 (1976). See generally 2 S. LADAS, PATENTS, TRADEMARKS, AND RELATED RIGHTS: NATIONAL AND INTERNATIONAL PROTECTION 949-62 (1975).

50 See S. LADAS, supra note 49, at 949.

51 See, e.g., id., at 949-50 (tracing evolution of petty patent principle from inventions of "tools, implements and objects of practical use" and stressing the now less valid proposition that utility models "must fulfill a technical utility function ... by means of a specific shape").

52 Cf. Reichman, Programs as Know-How, supra note 10, at 656-62.

53 See, e.g., Reichman, Designs and New Technologies, supra note 2, at 38. But see id. at 38-39 (discussing the United Kingdom's unregistered design right, enacted in 1988, which confers copyright-like protection on both functional and aesthetic designs).

54 See generally Reichman, "Legal Hybrids Between the Patent and Copyright Paradigms," work-in-progress first presented to the Forum on Intellectual Property Issues in Software, Computer Science and Telecommunications Board, National Research Council, National Academy of Sciences, Washington, D.C., November 30, 1989 [hereinafter Legal Hybrids].

55 See, e.g., Reichman, Designs and New Technologies, supra note 2, at 38 (citing authorities but also noting modern trend to lower the threshold of eligibility for the protection of functional designs generally).

56 See, e.g., 2 S. LADAS, supra note 49, at 952-54.

57 Id., at 955.

58 Id., at 952-53.

59 Id. For recent amendments to Germany's utility model law that lower the threshold standard to inventive activity from that of an inventive step, and that now embrace electronic circuit designs, see Hausser, Utility Models: The Experience of the Federal Republic of Germany, 26 INDUS. PROP. 314 (1987).

60 See, e.g., W. R. CORNISH, INTELLECTUAL PROPERTY: PATENTS, COPYRIGHTS, TRADE MARKS AND ALLIED RIGHTS 199-201 (2d ed. 1989).

[61] "Most patent systems ... have given the patentee no right to control the use or resale of goods which he has placed on the domestic market or has allowed a licensee to market there." W. R. CORNISH, supra note 60, at 200.

[62] See, e.g., F. PERRET, L'AUTONOMIE DU REGIME DE PROTECTION DES DESSINS ET MODELES 188-233 (1974); Pérot-Morel, L'ambiguité du concept de modèle d'utilité, in ETUDES EN L'HONNEUR DE R. FRANCESCHELLI 425 (1983).

[63] See generally Reichman, Programs as Know-How, supra note 10, at 648-662 (citing Galbi, Magnin, Kronz, and Kingston).

[64] See supra notes 50-59 and accompanying text.

[65] See Reichman, Programs as Know-How, supra note 10, at 662-67 (criticizing "a patchwork quilt of legal devices, complemented by an increasingly supple law of trade secrets, that has strained the classical system of intellectual property law to the breaking point").

[66] See, e.g., id., at 648-667, 714-17; Reichman, Designs and New Technologies, supra note 2, at 123-53; see generally Reichman, Legal Hybrids, supra note 54.

[67] See Reichman, Designs and New technologies, supra note 2, at 123-36.

[68] See id., at 141-53.

[69] Id; see further Reichman, Legal Hybrids, supra note 54.

[70] See, e.g., Goldstein, General Report, supra note 15, at 445-46. For a different analytical approach that also conceptualizes the computer as a tool, see Sookman, Computer-Assisted Creation of Works Protected by Copyright, 5 INTELL. PROP. J. 165 (1990) (discussing Whitford report in the United Kingdom and provisions of that country's Copyright, Designs and Patent Act of 1988, which allocates rights in certain computer-generated works).

[71] See supra note 3 and accompanying text.

[72] Cf. Ginsburg, Creation and Commercial Value: Copyright Protection of Works of Information, 90 COLUMB. L. REV. 1865 (1990); see also L. Ray Patterson & Craig Joyce, Monopolizing the Law: The Scope of Copyright Protection for Law Reports and Statutory Compilations, 36 U.C.L.A. L. Rev. 719 (1989).

[73] See, e.g., 1 P. GOLDSTEIN, supra note 1, at 197 (citing authorities).

[74] See generally Reichman, Programs as Know-How, supra note 10, at 693, 693 n. 288 (clarifying and reinterpreting Baker v. Selden, 101 U.S. 99 (1879); see also Reichman, Realist's Approach, supra note 1, at 970-76 (comparing and contrasting the views of Professors Goldstein, Kaplan and Melville Nimmer).

[75] See, e.g., Nimmer, Bernacchi, & Frischling, supra note 10.

[76] See generally Reichman, Designs and New Technologies, supra note 2, at 730-33. ("The Two-Market Conundrum").

77 See, e.g., Samuelson, Allocating Rights, supra note 15, at 1205-09; see also Stern, Framing Prints, Giving the Mona Lisa a Moustache, Speeding Up Video Games, and Marketing Add-On Software: A Comment on the Mirage Case, 37 PAT., TRADEMARK, AND COPYRIGHT J. 305 (1989).

78 The toolmaker may, of course, enhance his own reward if the tool embodies innovations that qualify for protection under patent laws, utility model laws, the Semiconductor Chip Protection Act of 1984 and its progeny in foreign law, or even a functional design law like that recently enacted in the United Kingdom. See, e.g., W. R. CORNISH, supra note 60, at 384-91. Indeed, toolmakers could expect particular benefits from a know-how law that yielded more discriminating rewards from applications of computer science than existing intellectual property models.

79 See generally, Reichman, Designs and New Technologies, supra note 2, 12-19, 37-42, 81-123.

80 See, e.g., R. Davis, Knowledge-Based Systems, 231 SCIENCE 957 (1986); R. DAVIS, Knowledge-Based Systems: The View in 1986, in AI IN THE 1980S AND BEYOND--AN MIT SURVEY (W. E. C. Grimson & R. S. Patil eds. 1987); R. Schank, What is AI, Anyway? AI MAGAZINE 59 (Winter, 1987); M. Minsky, Why People Think Computers Can't (198_).

81 See, e.g., Butler, Can a Computer Be An Author? Copyright Aspects of Artificial Intelligence, 4 COM/ENT 707 (1982); Samuelson, Allocating Rights, supra note 15, at 1205-09.

82 See supra note 43.

83 See, e.g., INTELLECTUAL PROPERTY ISSUES IN SOFTWARE, supra note 5, at 89-90.

84 See, e.g., id., at 66.

85 Cf. Directive on Computer Programs, Council of the European Communities, December 1990 (recognizing abuse of copyrights principle in the preamble).

86 See, e.g., In re O'Farrell, 853 F.2d 894 (Fed.Cir. 1988); Johnson, Patent Protection for the Protein Products of Recombinant DNA, 4 HIGH TECH L. J. 249, 250-51, 260-63 (1990). Of special concern are the "second-generation" analog proteins that may vary from patented "first-generation" proteins by as little as a single amino acid, which creates difficulties in meeting the nonobviousness test of 35 U.S.C. Sec. 103 (1988). See id., at 260-63; see also Baker & Parker, Obviousness: An Increasing Problem for the Biotechnology Patent Applicant, 7 INTELL. PROP. BULL. 2 (1990). Other problems include the difficulties of obtaining process patents, see e.g., In re Durden, 763 F.2d 1406 (Fed. Cir. 1985), the long examination period, disclosure requirements, and the need for university professors to publish their research results. See, e.g., Eisenberg, Property Rights and the Norms of Science in Biotechnology Research, 97 YALE L. J. 177 (1989).

87 Cf., e.g., Hogle, Copyright for Innovative Biotechnological Research: An Attractive Alternative to Patent or Trade Secret Protection, 5 HIGH TECH L. J. 75, 76-77 (1990); Smith, Copyright Protection for the Intellectual Property Rights to Recombinant Deoxyribonucleic Acid: A Proposal, 19 ST. MARY'S L. J. 1083 (1988); Kayton, Copyright in Living Genetically Engineered Works, 50 GEO. WASH. L. REV. 191 (1982).

88 Cf. supra notes 50-54 and accompanying text.

89 See, e.g., Davis, supra note 37; Reichman, Programs as Know-How, supra note 10, 659-62, 696-98. The concepts in the text have benefitted greatly from unpublished discussions with Mitchell Kapor, coinventor of Lotus 1-2-3.

90 See generally Davis, supra note 37.

91 See most recently Reichman, Proprietary Rights, supra note 34; see also Samuelson, Innovation and Competition: Conflicts Over Intellectual Property Rights in New Technologies, 12 SCI., TECH. & HUM. VALUES 6, 14-16 (1987).

92 See, e.g., Butler, supra note 81, at 728-44; Cf. Spoor, Expert Systems and the Law--An Outline, 1990 Copyright 330, 333.

93 See Reichman, Designs and New Technologies, supra note 2, at 127-33, 141-53.

94 Cf. Eisenberg, supra note 86, at 190.

95 See, e.g., Newell, The Models Are Broken, the Models Are Broken, 47 U. PITT. L. REV. 1023 (1986).

96 Cf. W. KINGSTON, INNOVATION, CREATIVITY AND LAW 83-85 (1990); see also E. MACKAAY, ECONOMICS OF INFORMATION AND LAW 115-17 (1982).

97 See, e.g., Lucas, General Report--Data Bases and Copyright, in L'INFORMATIQUE ET LE DROIT DE L'AUTEUR 332, 333-36 (proceedings of the ALAI Annual Congress, Quebec, Canada, Sept. 26-30, 1989); Spoor, Expert Systems and Copyright, supra note 14, at 103-04.

98 See generally Reichman, Legal Hybrids, supra note 54.

COMPUTER-PRODUCED CREATIONS

by

Robert Barr
Partner, Townsend and Townsend
Palo Alto, California
United States of America

I am a practicing patent attorney here in Palo Alto. I would like to tell you about a recent visit that I had from a new client. I have that client's permission to tell you a bit about his product. I will refer to the client as Mr. Dekker, which is not his real name. The product is a code generator. He calls it his seventh generation code generator. Some of you may remember the fourth generation languages that were going to eliminate programmers--apparently they did not succeed. This seventh generation product, in response to functional specifications from the user, generates computer programs. Those computer programs are original in two senses. Firstly, they are the type of work that we would normally consider copyrightable if created by a human author. Secondly, they are also original in a more unique sense, in that they do not contain any code from the generator itself. Previous code generators contain sub-routines or signals of code that are strung together to form the output program. But Mr. Dekker's code generator actually writes new code.

I talked to Mr. Dekker, and we are planning to patent some of the inventions contained in the program, in addition, of course, to securing copyright protection. The question quickly came up as to who owns the output of this program. In Mr. Dekker's view, this is where the real value is, and he thinks he has created a tool that will create commercially useful programs. People may have to add to them and adapt them, but the generated code is source code, and it works. According to Mr. Dekker, the programs are "bug-free." He asked me about authorship of the output, and I told him that I have been researching that issue for a conference. It is a very difficult question. He said "what do you mean it is a difficult question? If I am the author of the source code in this program, that means I own the source code; according to the copyright law, that means I own the object code too. In fact, I own the screen displays, the user interface, and from what I have been reading lately, I own the non-literal elements of the user interface. It is all output of my program--why is this different?" I told him that the compiler is an example of a program whose author does not own its output. There is no intervening human authorship when a compiler generates object code from source code.

Mr. Dekker was rather indignant about comparison to a compiler because this program is much more than a compiler--for one thing, the user only inputs functional specifications, but the output is procedural code. As he said to me, "this functional specification is purely in the realm of 'ideas' even under the broadest definitions of Whelan v. Jaslow. The program takes in ideas and functions and puts out expression." He said, "what about derivative works?" I said, "that's interesting because if, in fact, the program had incorporated in its output some of the code from the generator, I would treat it as a derivative work, but you tell me that is not the case. Maybe we should consider the generator itself as the creator of the program. The problem in the United States and other countries is that there are statutory and constitutional obstacles to calling the computer an author."

His comment was, "you attorneys had better get ready for the twenty first century. Haven't you heard of the Turing Test? You will not be able to tell from these programs whether they were written by a human author or a computer author, and it is very important that they be protected and owned by somebody."

Then he thought about it for a while and he said, "maybe that is not the issue at all, maybe ownership is just a construct of lawyers anyway. I am in business to make money." He proposed a licensing program to me. He figured if the copyright law and the patent law are not clear, maybe what we need to do is put together a licensing program that will allow him to recoup his investment and make a profit. He had several interesting ideas, and I will share four of them with you and the legal issues they raise.

Number one: he wanted royalties on the output, per unit royalties on each copy of the program that is sold. He figures that way he and his customers can share the risks and the benefits of using the generator. Whether these programs are actually marketable remains to be seen, and the royalty structure is classic for sharing the risk between the licensor and the licensee, so why not royalties on a per copy basis? Number two, as a given program starts selling a lot, after the first thousand or hundred thousand copies, he wants to raise the royalties. Sometimes you will see it go the other way--in this case we want to share the risk but we also want to share the rewards, so after a certain number of copies we want to raise the royalties that he is going to receive on the output of the generator.

Number three was very interesting. He had sure been paying attention. He wanted a grant-back. Not a grant-back from his licensees of ownership rights in the generated programs, but a grant-back of non-exclusive rights under any patents and copyrights in the programs produced by the generator. He wanted the right to sub license to others the right to use, reproduce, and distribute these programs. The reason for that is interesting because, while the program will produce unique output, he envisions that his licensees may end up suing each other. His programs might have similar structure, sequence and organization, maybe similar look and feel, and the last thing he needs is his licensees suing each other because, as he said, "they may tell you that I do not own the output, but they are going to look to me when they get sued." Again, maybe ownership is not the whole issue, but they are certainly going to turn around and look to him if they get sued, so by having this grant back he could avoid the problem of his licensees stepping on each other.

Fourthly--again so that his licensees do not interfere with each other--he proposed certain field-of-use restrictions on the output. The licensees will use the program, some companies will be able to use it to generate application programs of a given type, others will be restricted to internal use of the output--some big companies can use this like the fourth generation programs were used to create in-house programs, to do away with programmers. And still further, he wanted to reserve to himself and his company the right to create artificial intelligence programs. Mr. Dekker had concluded that we will solve all the problems of computer generated output by contract, good old private law. That sounds good, but what about Lasercomb? Lasercomb is an interesting recent case applying the principle of copyright misuse to a software licensing program to prevent a licensor from suing for infringement. It is a doctrine which is very well founded in patent law, and while it seems to have faded in importance for various reasons, it is still raised in many cases. It is not the same as anti-trust law, but it is related to anti-trust law. Misuse is a sort of an unclean hands defense, that an accused infringer can raise against enforcement of a copyright, even if the party who is being sued for infringement--as was the case in Lasercomb and in most of these cases--is not a direct "victim" of the misuse. The accused infringer says you cannot enforce your copyright because what you have been doing with it is against public policy, it is anti-competitive, and it extends the scope of your copyright beyond its lawful realm. In patent law I think the phase is "to expand the scope of the patent monopoly to unpatented goods," so you can see it might have some application here.

When I went to research this issue, I quickly found there is not much in the computer law books or the copyright law books about copyright misuse. Copyright has already been stretched to cover functional works and you will not find much in the way of distinction between output and a copyrighted work, because maybe a copyrighted work is not supposed to have output. For the last fifteen years there have been programs with output, but there is not a lot of law on that distinction, whereas in patent law the distinction is an old one. An invention of a process or a machine is usually a distinct invention in this country, from the product produced by the process or machine. And it turns out that there is a lot of misuse law and a lot of anti-trust law regarding what a licensor can do with his intellectual property, whether he can extend that to restrictions on the output as proposed here.

I am going to briefly cover what I discovered in this research. I do not know how purely American it is, I think there is a related principle in British law--the principle of non-derogation of grants. If there are no such considerations in your country, you may consider this a creative licensing program and maybe we can deal with the issue of ownership of computer generated works, while we are waiting for other solutions, by license arrangements. But here we have to deal with the issues of copyright misuse and potentially anti-trust so I will cover briefly what the problem are in American law.

As far as the per-unit royalties on the output, I am not sure the misuse doctrine applies here. There are patent infringement cases where reasonable royalties were used as a remedy by the courts based on what a willing licensor and licensee would have agreed on. There are cases like the General Motors v. Devex case, where millions of dollars were at stake on a process for making bumpers. The bumpers were not patented--the process was. And when General Motors used the process without permission they had to pay a per-unit royalty on bumpers. There are other cases involving per-unit royalties on water heaters and other unpatented articles made by infringing a patented process, so presumably the courts did not see a misuse or anti-trust problem.

There are royalty cases, misuse cases, on the issue of whether your royalty base can cover only patented products. It is common in some industries to charge a royalty, not only on a patented product but on all of the products of that type sold by that customer. There are issues as to whether that is coercion or condition of the license, but I do not believe there is really any problem with this royalty here. There are practical enforcement issues, obviously, but Mr. Dekker does not seem concerned about that.

Number two: a higher royalty as the volume goes up is usually treated under misuse law as an output restriction, such as in the Q-Tip case. There was a famous case in 1951 in this country where the Q-Tip people sued Johnson and Johnson because they owned a patent on a machine for putting cotton on both ends of the cotton swab. Johnson and Johnson raised the misuse defense.

They were not a licensee but it turned out that the Q-Tip people had limited the number of double sided cotton swabs that could be produced by its licensees. So it is exactly the kind of test we are talking about here in a patent context. The court held in that case it was not a misuse, but other cases have gone the other way and left the door wide open to examination of the reasonableness of such restrictions. The criteria is basically whether they are ancillary to the lawful purpose of exploiting the patent, and are subject to misuse and anti-trust considerations.

Before I cover the last two points in the license program, I want to point out an obvious difference between a Q-Tip and the output of the generator. Maybe it is not obvious. A Q-Tip is not intellectual property. The machine that makes the Q-Tip must generate thousands of Q-Tips. The code generator produces intellectual property. It will generate one copy of the code and you will use a copy utility or something else to generate multiple copies. Whether or not these misuse cases apply, what is really unique is that intellectual property might be produced by these artificial intelligence tools. Also, when you have patents, the "use" right is really the basis for allowing these kind of licenses. The patent holder controls the use of the machine and therefore can, in some circumstances, put restrictions on how it is used. If we only have a copyright for this generator, there is no use right under American copyright law as such. We license use in all our software licenses, but it is not one of the exclusive rights.

The grant-back issue, number three, has also been the subject of a lot of misuse and anti-trust litigation, and it is not something that you want to take lightly. One of my early reactions to the uncertainty in the law of computer-generated output was to have the licensee grant back all of the rights. You could run into trouble with that in the United States under these principles. It is clear that there are many factors that the courts consider, such as the market position, market dominance or lack of say by the licensee. For example, it was not OK for GE to demand an assignment of improvement patents from its licensees on light bulbs, but it has been held OK in other cases.

The last issue raised by this licensing program is the field-of-use restrictions. Interestingly, in the mid-70's in this country the Justice Department created a list of "no no's" and on the list, for anti-trust purposes, was that restrictions on the sale of the unpatented product of a patented process were unlawful extensions of the patent monopoly. That is not their position today, but they did bring suit in a 1981 case against a company that licensed its machine in a way that one licensee was allowed to distribute the output and the other licensees were not. Only one had the exclusive license, and under this principle the Justice Department brought an anti-trust suit and lost.

As we look at the need to balance the incentives, everybody seems to agree you do not need to give incentives to the machine, the computer, certainly not financial incentives. We do need incentives for the creator of the artificial intelligence program, we do need incentives for the user to put the output into commercial circulation when appropriate. I think the licensing aspects illustrate the way you can balance those incentives, to the extent the market will bear and to the extent that we and the legal system will decide that it is or is not anti-competitive.

COPYRIGHT PROTECTION OF COMPUTER-PRODUCED CREATIONS

by

Péter Gyertyánfy
Hungarian Bureau for the Protection
of Authors' Rights (ARTISJUS)
Budapest
Hungary

Works Created by Computer

1. The human mind cannot resist exploring new phenomena as subject matter or means of artistic creation. It is hardly by chance that the first "computer programs," created about 160 years ago by Lady Lovelace, the daughter of the famous English poet Lord Byron, were intended for composing works, namely musical works. The computer has an impact on nearly all fields of artistic creation as it has on nearly all other human activity. The computer is used as a means of or an aid to the creation of picture sequences, films, musical works, sculptures, translations, literary works, poems, etc. Computer-aided design and image animation have become elements of the modern creative processes.

2. For the purposes of this paper, computer-produced creations are products achieved by or with the aid of computers which correspond to the criteria of traditional creations in the literary and artistic domain. The only difference is to be found in the use of computer programs, in the process of creation. However, the use of technical means in the artistic creative process is neither new nor alien to copyright protection (see, e.g., photography, film production). The question is whether computer works represent a new quality in this respect.

Should Computer-Produced Works be Subject to Legal Protection?

3. The abundance of computer works and the complexity of the technical aspects of their creation give the impression that we are faced with an entirely new situation. The automatic protection of "works" by programs could limit the interest of producers in their legal protection. They might prefer not to claim any rights in the results of running their work-creating programs in order to facilitate distribution. The short-term interests of consumers correspond to that attitude.

4. However, as experience shows, automatic exclusive rights are needed by the producer-creators of video games and other computer works, just as for traditional works and computer programs. The moving forces are the interests of the creators for whom the basis of material compensation is the exclusive rights and the producers' interests in the security of investment, i.e., in a property-like protection excluding duplicate products of others from the market. The long-term interests of consumers are in continuous access to choice of works.

What Kind of Legal Protection is Applicable?

5. If we emphasize the common features of computer-produced creations and traditional works, it is evident that they should be governed by the general principles of copyright protection. This is all the more so if we think of the necessity of prompt international protection. Under "copyright" I include, of course, not only the legal institution of authors' rights but also those rules which grant exclusive rights to the producers of phonograms, and of products not falling under the notion of "work" (e.g., sequences of images, "Laufbildschutz" in the German law) be they protected under the name "neighboring rights" or within the body of copyright legislation.

6. But, is copyright an applicable and adequate legal institution for computer-produced creations? At first blush, we could say, in the words of the WIPO Expert Committees of 1980-82, that these creations are eligible for copyright protection if they satisfy the requirements of copyrightability, i.e., if they have a sufficient degree of originality and they are the result of a creative effort. Value, mode or form of expression, purpose and destination are irrelevant; ideas and practical results are not copyrightable.

7. This is, of course, a general approach. The practical conditions of the "identifiable, expressed form of an idea," which is a work, may differ from one national legislation to another. Our answer to the above question regarding a specific work or case cannot be precise unless based on the national legislation in question.

8. German court decisions have produced, e.g., a specific requirement for the "work quality" ("Gestaltungshöhe"), for the level of individual creativity regarding computer programs. This has not been without consequences for the copyright eligibility of video games, either. For the copyrightability of video games, the level of individual creativity in the computer programs defining and controlling the video games must be above the everyday average production of that field--according to certain German high courts.

9. Other German high courts maintain, however, that there is enough creative specificity in video games, which are similar to films in general; all the effects of the interventions of the video player are pre-programmed; the individual "work capacity" is given not only in the case of an expression of thought reflecting the personality of an individual but also if it reflects, as content, a thought resulting from an individual mental activity.

10. Coming back to a more general level of deliberations concerning eligibility for copyright, we have to turn to the identified difference between "traditional works" and "computer-produced creations": to the use of computer programs. Before now, the process of creation has been irrelevant from the point of view of copyright eligibility of the result. The difference between mechanical or handicraft production was relevant for copyright eligibility of the works of applied art only in certain national laws, and many decades ago.

11. One of the basic ideas of copyright law is that the possible results of the work, i.e., the message content, the ideas or problem solutions, are outside the scope of protection. This principle has been proved and generalized on different traditional work categories. Computer-produced creations are wholly or partly a result of another protected "work" of the

computer program. The program is, usually, an identifiable original expression of thought, describing--individually--a trail of thought of mathematics, in other words, a route to a problem solution. In our case, the problem solution, the result of the program-work, is the computer-produced creation. Can it be protected despite the above general principle?

12. The key to the answer can be found in the double nature of computer-produced creations. It is, on the one hand, the result, at least partly, of the functioning of the algorithm; in other words, it is the solution of a mathematical problem; on the other hand, at the same time it can also be an individual expression of thought. If the signs of originality of its inner and outer forms are identifiable, if it corresponds to the characteristics of the works created/produced in a traditional way, a computer-produced creation can be eligible for copyright protection. The conditions of eligibility are to be considered in regard to the "creation" separately, regardless of the process having led to its birth.

13. As a matter of fact, we have also to state precisely the general principle which was the starting point of the above deliberations to a certain extent: copyright protects the content and result--as an inner form--even in the case of certain works created traditionally (translation, plot and characters or novels, etc.).

Specific Questions to be Answered

Multiple Works

14. This problem emerges first of all in the field of "automatic creation" (the computer functions as a "robot author"). Namely, the more latitude the user of the computer has for the intervention into the running of the program and for using or not using its suggestions, the greater the probability of the birth of distinct works (e.g., computer-aided design). In the most simple of possible cases, the program is capable of reproducing one work only. The result is 100 percent foreseen and fixed in the computer program. If the running is repeated, we can produce multiple copies of the same work.

15. Programmed creation, however, may also produce different versions if it is based on random production and selection of items as a first step of the program. Assembling and organizing rules (e.g., music-composition rules) are being applied to ever newer "raw material." In this situation, I think, all the existing versions have to be judged separately for copyright eligibility. Nevertheless, the programmer must foresee the various possible versions in their decisive major characteristics at least as a precondition of the "work quality," of copyright eligibility.

16. Legal literature draws attention to a possible danger in respect of future multiple works. Theoretically, the rules of a work-creating program can be defined so broadly that the running of the program could result, in theory, in all possible works of a given work category. The program would cover an indefinite range of works. Could the author of the program claim authorship in all these future works, thus excluding others from creation in this field? I doubt that this difficulty is a threat, even in theory. Namely, parallel creation, unintentional "plagiarism" is not an infringement

in any copyright legislation. "Parallel works" can exist. Further on, you cannot grant copyright to a number of possible future works of art. This would be a patent approach putting under exclusive protection the future results of a patented method.

The Problem of the Work Category

17. It has been emphasized that computer-produced creations should be qualified according to their own features for the purpose of copyright eligibility irrespective of the process used to create them. This also leads unavoidably to a general conclusion regarding the genre of such works: they are musical works, literary works, graphic or audiovisual works, etc., according to their main features which correspond to the characteristics of these traditional categories of works.

18. However, we can also argue that since the creative elements incorporated into the computer program which assist or determine the produced work are transmitted into the latter, the computer-produced work is of the same genre as the computer program.

19. Computer programs as works are defined differently in the above respect under national legislations. According to the recent EEC Directive on Computer Programs, the "software work" is a "literary work." This would mean that the genre-specific rules of the given national copyright legislation concerning the computer program (in the case of the EEC: rules on literary works) are to be applied to every kind of computer-produced creation. These genre-specific rules may include wide areas from the specific originality requirement through the contracts for use to the exceptions (fair use, etc.). At this point, I have to remark that the EEC Directive is self-contradictory when it qualifies, on the one hand, the software work as a "literary work" and, on the other, prescribes a specific exhaustive catalogue of exceptions, alien to existing exceptions restricting literary rights.

20. In my view, the characteristic and the effects of the computer program assisting or guiding the creation of a computer-produced work should not be overestimated. The separately judged work qualities of computer-produced creation--contrary to the case of data base works--are more decisive for the purpose of genre-qualification. Otherwise, we could run into the practical difficulties described above.

Can a Computer be the Author of a Work?

21. The generally accepted answer is, of course, no. The computer and the computer program used for creation should be considered as technical means used in the process of creation for achieving the results desired by human beings. The author of the computer work is basically the person who produced the creative element in the computer-produced work and the program guiding the creative process. But, as we know, the devil is in the details. There is a rather wide range of further questions open to interpretation even if we put aside the non-specific problem of copyright ownership in commissioned works and works made by employed persons.

22. In the theoretical case where a program is capable of producing one work only, it is clear that the rules of creation are incorporated in the computer program. Therefore, the only person who can be regarded as the author of the work in question is the programmer. The possible creative cooperation of further persons contributing to the program (technicians, date base input engineers, experts, etc.) is a non-specific question here. Finally, original authorship is not subject to agreements, and the contracts of the collaborators may have a practical relevance in this respect also, at least until the contrary of "agreed" authorship is proved.

23. The answer is still obvious to our basic question when the result of an automatic computerized creation is a wide range of "generated" variations of musical works, designs or poems and it is the program writer who makes his choice in designating one of them as "the work." The development of software rules and the recognition of the result as a work together qualify the programmer as "the author."

24. On the contrary, should the person making the choice between the output versions be some other than the original programmer, his authorship or co-authorship should be denied. To single out, to find, to choose an item from various existing ones cannot make someone an author, a creator. While this act can be a final, inherent step in the process of creation--the artist may try his hand at different rough sketches and variations--this alone can not amount to "creation." As a matter of fact, I doubt whether the series of output variations can research "work level" in such a situation in practice.

25. Identification of the author of the computer-produced work is, of course, the core of our topic, and we encounter the most difficulties in the field of the "computer-assisted creation." Such creations come from situations where the computer program does not seem to contribute any creative element to the work, but serves as a tool only. Such are many of the CAD programs. The author of the computer-produced work is, in this case without any doubt, the user-architect, the user-designer and so on.

26. Certain programs are designed for very general use, e.g., for the composition of music in a fairly wide range of styles. The artistic form of the result and some features of the works to be produced are influenced by this kind of program, but not fully. The result is not prejudged. I think that authorship of the works produced with the aid of such programs should be attributed to the user. It goes without saying that the originality in, and work quality of, the creations is a further question separate from that of authorship. It was discussed in this paper earlier.

27. The highest developed level of computer-assisted creation is where the user can alter the course of the programs and engage in a real two-way dialogue with the computer. Such programs make creative contributions to the final result but leave space for the user to add his or her artistic contribution. Unfinished, incompletely defined musical compositions (matrix works) can be finalized by the user. The computer can provide visual models, planned to be executed or modified. Here we are dealing with more than one author: with the program-writer and with the user-co-author.

28. Consequently, we can say that the human author is never absent if there is a computer-produced work. Admittedly, it is sometimes difficult to identify him/her. However, it is a technical, not a theoretical problem and the answer should be the application of the general rules of co-authorship. The allocation of original authorship to others than the human creators is no more justified here than in the case of traditional works.

29. It is undeniable that such co-authorship may cause certain problems. Just to mention one consequence: the heirs of the co-authors of a joint work enjoy a special term of protection regime.

30. As to the exercise of rights, we can assume that the licenses for use and distribution of work-producing programs include the right to use the programs for their very purpose, for the finalization of the inherent works. Practical problems begin with the necessary authorization of further use of such computer-produced works. Acts subject to the rights holders who have control over computer-produced works are the same as in the case of traditional works. National copyright rules regarding each work of art category and specific use apply.

Preexisting Works Affected

31. From the point of view of the law, the use of preexisting literary or artistic works (created by traditional methods) for the purpose of a computer-created work is a separate dimension, distinct from the area of original computer-created works discussed above. Data processing methods are often used to adapt well-known classical, public domain musical works into fashionable pop-music. This can, of course, also happen with protected works of the same genre. Certain graphic programs can produce gradual modifications or slightly repeated patterns of an original model, or a given picture. Computer-related literary creation is less developed than the musical or graphic field. This is probably due to difficulties still remaining in giving semantic abilities to the computer. The computer is used in this territory of art, therefore, mainly for translation or modification of preexisting works; the computer can also build or re-structure new literary works form the word content of stored works.

32. Under the international conventions and national legislations, adaptations enjoy a separate, but conditional copyright protection: the adaptation may not prejudice the copyright in the preexisting work. The exclusive right to authorize adaptation of a work is granted in different ways (moral right, economic right, or both as a moral and an economic right) under differing legal systems. It also has its limits: the right to make certain changes obviously necessary for authorized use.

33. All these rules and the problems of their application to computer programs are well known. Here I have to add two specific aspects only. The requirements for a derivative work--the degree of the necessary creative contribution--are different in the different national laws and even more so in practice. These apply, mutatis mutandis, to computer-reproduced works also. The arrangement and adaptation of musical works in the public domain is a sensitive issue. General aesthetic, cultural as well as economic objections have been raised against such a "misuse" of the valuable cultural heritage in classical-romantic repertoire. Computerized mass-adaptors have to bear in mind that there are different legal means for the protection of the integrity or works in the public domain in many national copyright laws, even apart from the copyright legislation.

COMPUTER-PRODUCED CREATIONS

by

Akin Thomas
Managing Director
Heinemann Educational Books
Ibadan, Nigeria

Distinguished members of the intellectual property community, this last paper of day 2 is maybe the first, and possibly your only opportunity to take a break from serious discussions of engineering issues and sensitive legal points. Between yesterday morning and this afternoon you have listened to many possible definitions and restatements of the crucial issues from all possible major angles. All that is now left are the minor angles that we all tend to easily overlook in a forum like this. These minor angles revolve around the implications in the area of copyright for computer creations. I would like also to discuss the implications for developing countries. But first, I would like to thank Dr. Arpad Bogsch for inviting me to present a paper on computer produced creations at this seminar.

I come from a country where the dawn has only just broken in terms of computer technology and the mass use of computers in the fields of commerce and industry, education, banking, and information dissemination in general. Therefore, the opportunity to examine the copyright problems of computer systems has been a rewarding one for me, and I would like to share my thoughts with you in a general way, mostly because I have no specialist opinions in the specific areas of computers and new technologies beyond my familiar field of book publishing.

Computers have the means of collection, storage, organization and dissemination of information by electronic means, to a restricted specialist audience or a general audience via satellites, cables, or telephone lines. In two areas, computers are of immediate relevance to a book publisher. First, equipment with sophisticated word processing capabilities offers the capacity for storing written materials electronically, in machine readable form on floppy disk or optical disk, which can be used for reproduction via on-line or off-line laser printers. The second area is the data base or data bank, which offers limitless prospects of access to information in a well organized machine readable format. This information can be made accessible to authorized users either on-line via cables or off-line via magnetic disks like CD-ROMS, which can read as if on the screen. If required, the information can be transmitted to another compatible computer unit or reproduced on floppy disks or optical disks. The disk, or other computer software available to the owner of intellectual property, offers the owner of such intellectual property the opportunity to prepare his work in machine readable form only, making it possible for him to store such information safely in a computer memory or on disk.

One asks--what is or should be the copyright position with respect to products of computer systems? Because materials stored in computer programs are usually protected works, they enjoy, for example in Nigerian copyright law, the same protection as all intellectual property in fixed form. This protection covers the arts of storage and retrieval from computer memories, from data bases and from electronic main libraries.

The rights protected include:

1. the right to make or authorize the making of translations, adaptations or derivative works;

2. the right to reproduce any work;

3. the right to make the work available to the public by direct communication;

4. moral rights.

My view is that in relation to the general principle of protection of intellectual property, the protection of such material is desirable. There is also a need to examine the implications from the point of view of developing countries especially in terms of the use of computer systems for access to all the creation of protected works and the users of protected works in order to stimulate creativity of authors and not hamper the dissemination of works by means of computer technology. Developing countries can only hope, as far as development through education and books are concerned, that computer technology--in an age in which non-book materials are fast replacing books--will not mean the death of publishing and knowledge in the developing countries.

Bearing in mind the obvious fact that the developing countries have neither the technology nor economic resources to cope with new and ever changing development, it is also possible that intellectual property now vulnerable to piracy in book form may be better stored in safer computer terminals and CD-ROMS in the future. There is already a trend toward making new scientific and technological research available only or mainly through electronic outlets, which most developing countries cannot afford and therefore have no access to. Unfortunately, as far as development by education and books are concerned, recourse to the traditional printed book will offer little hope in new professional and technological fields, since most new developments may not find their way into books for a long time, if ever, in a nation in which every home in the developed world can afford the personal computer and have access to electronically published materials. But developing countries are pushed further and further away from such products. It is also just possible that as a book becomes less relevant to the developed world, all developments and improvements in traditional book printing as it is known today, may cease.

The experience in the electronic media is worth noting. Very few developing countries, including Nigeria, can afford to develop television or radio as means of education, information and entertainment, because of the cost of equipment and the fast changing technology. Such countries have therefore become less and less able to disseminate information, while developed countries have established cable television networks to fill the gap. Unfortunately the intention and orientation of such developed countries are not congruent to the national interests of developing countries. Copyright culture is getting stronger, however, and the developing countries will be paying through their noses for cable television services for decades to come. The implied information/colonization trap is very worrying, as far as the international implications for the poor developing countries.

The point of the last few paragraphs is not that computer technology or computer systems are evolved purely for developed countries, but rather to draw attention to the fact that unless developing countries are assisted out of their present state of helplessness in educational and cultural areas through technology required to sustain their basic needs of educational enlightenment and cultural mobilization, their future is gloomy against the back drop of recent developments in the electronic publishing and computer fields. Such a prospect does not augur for peace and friendship in the international community and the protection of intellectual property rights at the global level.

We may therefore wish to remind ourselves of the principles and noble intentions of copyright and the history of the copyright system as a social and economic tool. As we all know, the copyright system has developed over the years as a means of protecting creative people and their works. Copyright not only provides creative persons with legal protection and rights, by the exercise of which they may earn a living, but also serves to encourage the development and propagation of creative works by giving to such creative persons and authors the right to exercise certain rights. The unauthorized reproduction and distribution of such works would be so easily accomplished to the detriment of growth in creative activity.

It may be said that the purpose of copyright laws is to spell out the conditions under which copyright material can be used or not used. And since copyright itself is not an unlimited right, conditions exist under which the owner of copyright cannot prevent the use of his copyright property. At the international level, the United Nations General Assembly in its 1948 Universal Declaration of Human Rights acknowledges the individual's rights of access to works of literature, education, science, technology and art. The Declaration states in Article 27, paragraph 1 that "everyone has the right freely to participate in the cultural life of the community and to enjoy the arts and to share in the scientific advancement and its benefits". This same UN body shows that it does not overlook the crucial and legitimate interests of those who create works of culture, science and art as is evidenced Article 27, paragraph 2 of by the Declaration, which states that "everyone has the right to protection of the moral and material interests resulting from any scientific, literary or artistic production of which he is the author".

The creator's rights are generally protected by national copyright laws with the obvious limitations of jurisdiction restricted to national boundaries. However, in order to ensure the protection of copyright material beyond such national boundaries--because most creative works, especially computer materials, travel across national boundaries at will--it is important for such intellectual property to enjoy legal protection across national boundaries in such a way as to ensure ideally global protection. In no other area is this more crucial than in the fast growing field of new technologies, especially computer programs which present a special problem in terms of copyright interpretation and administration.

And Stanford University, because of a special position in this area, looks like the most ideal place at this point in time to tackle the problems of the copyright implications of computer programs and artifical intelligence generally. Some earlier papers at this seminar were devoted to explaining this scope and background including the main categories of artificial intelligence, but for the sake of completeness I am tempted to provide some definition. The definitions are stated in the paper--artificial intelligence, the expert system, and computer-assisted works, I won't bore you with all that. At this point it is probably useful, however, to remind ourselves that copyright has its root in the acknowledged need to compensate human--<u>human</u>, not mechanical--creative efforts, that is, the intellectual activities performed by human beings, and thereby the encouragement and the continued development of creativity. The products of creative endeavors protected by copyright are things which can be easily seen and/or quantified like books in published form, works of art, dramatic works on stage and theatres, musical works performed at concerts, etc. However, artificial intelligence attempts to extend the products of creative endeavors to some products of non-human elements created by man. We have examples of this in the output of a digital sampler in the creation of what is called original music without the efforts of a human composer, or computer-created original paintings, and very soon we may be having computer-created works of literature.

Already in the area of warfare, as we have found out from the recent Gulf experience, the computer made so many things so much easier that the "Mother of all Battles" almost did not get fought at all. In all these achievements of the computer there is an implied danger that the computer and computer-produced creation may kill creativity in human beings--computers by their very nature will always be able to do many things better than humans, or at least faster.

Let me not digress too far, and return instead to the issue of artificial intelligence and copyright. The multiple creativity involved, that is, acknowledging the original efforts of the designer of the program, and the obvious mechanical achievments of creation of the computer, and the fact that the efforts at some point cease to be human efforts, raise doubts about the place of such systems and intellectual products in copyright. Should they be classified under copyright or as inventions and/or industrial property? We must remember that intellectual property derives from property rights of intellectual creators, which is what copyright is all about.

Property rights, of course, relate to movable and immovable property--a car and a house are examples of the respective categories over which the owner has an exclusive right. While intellectual property relates to creative literary and artistic works and inventions, and although Article 2 of the Berne Convention lists the literary and artistic works protected by the Convention, the items are mainly in the literary, scientific and artistic domain, in varied modes of expression like books, musical works, choreographic works, drawings, paintings, architecture, sculpture, engraving, etc.

However, the original modes of expression have been enlarged to include in the area of scientific inventions the computer program, which involves intellectual or mental steps and instructions controlling the operation of a computer designed to perform a task. The program is produced by man's intellect, but understood and interpreted by the machine or the computer rather than by human beings. If we look at computers as purely mechanical products which are themselves subject to patent, can we regard their products or creation as anything but patent related? I think the answer would have to be yes, if a computer program is involved, because such program is an intellectual work and its author is the designer of the program. The output or creation of the program is therefore, by extension, the creation of the intellectual factors put together by the designer of the program.

The products are the computer assisted or the computer generated works of an author, with the mechanical assistance of a computer. However, in the case of computer-assisted programs where only minimal tasks are incorporated, one can say that the computer program does not complete a story—it only provides a variety of plots to choose from and leaves it to the user to weave the plot into a story. The user in such a case is the owner of the copyright, and the computer or the computer program remains only a tool.

The joint Unesco/WIPO meetings of 1980 and 1982 did quite a bit to lay the ground work for discussion and resolution of the problems. Yet, after these meetings and follow-up meetings between then and 1989, it does not appear that we have arrived at a full consensus. I hope this Stanford seminar will get us to a conclusive point on this matter.

Let me address some of the issues of computer creations from the point of view of a developing country. The first point relates to sovereignty and national economic interests, the second point relates to the effects of computer and satellite technology on cultural and educational development of developing countries as a result of satellite technology, and the third point is an example of how the computer or a computer-produced creation can kill creativity. Here I want to cite the example of something I ran into in London two weeks ago. The talking drum is an old and familiar article. I was horrified to find a talking drum in fibreglass format complete with synthetic strings. A brochure claimed that this product is better than the original talking drum. When it is patented, we will be expected to pay for the copies of this computer-produced creation. Very soon we will get human beings created by computers, robots with flesh and some blood in plastic veins. If there are any such humans within hearing of this auditorium I ask their forgiveness.

I would like to move on to the end of this and touch the question of distribution rights, especifically for books. The author should have a right to state, for books and the computer, which mode of expression is being authorized for exploitation—is it a sales mode or is it a rental mode. If it is sales then rental should be prevented, and if it's rental then sales should be prevented. It means two contracts, but I see no reason why this should not be done. The usual practice in publishing would have been to lift the heavy investment onto the rental mode like the hardback copy of books which are usually kept in libraries. I think this is an option that is worth exploring, if the author is going to benefit from the multiple user implication of the rental mode.

There are a few other points which I would have loved to draw your attention to, but I would like to end by saying that the copyright culture should grow on the principle of adequate compensation for creators. The copyright culture should concern itself with compensation if it is to result in a propagation of culture, information and education. That includes ensuring it contributes to the availability and affordability of such works in our countries, developed or developing. New technologies may expand the economic potential of developed countries in the areas of entertainment, information and education. But if this expansion is not matched by an obligation to examine the needs of the global market and the audience that copyright protects, then we might as well say good-bye to the good way of developing countries in copyright matters.

COMPUTERS AND AUTHORSHIP: THE COPYRIGHTABILITY OF COMPUTER-GENERATED WORKS

by

Arthur R. Miller
Bruce Bromley Professor of Law
Harvard Law School
Cambridge, Massachusetts
United States of America

I. Introduction

We all have heard about the proverbial roomful of monkeys striking the keys of typewriters (or, in this day and age, word processors), with one of the monkeys eventually coming up with Shakespeare's "Hamlet." Those of us who teach copyright law often put the question to our students (but do not answer it): How should the copyright law treat the work? Aside from the question of independent creation, the most interesting inquiry is whether the monkey's contribution constitutes sufficient "authorship" to make the simian "Hamlet" a copyrightable work.

We are not likely to encounter our hypothetical Shakespeare-like monkey in the real world,[1] but some people believe that the non-human author is already with us--that thanks to advances in artificial intelligence, computers or computer programs (or both, through their joint efforts)[2] now are producing works whose origin cannot be traced directly to a human author.

They believe that, although our simian Shakespeare may be purely a classroom construct, a cybernetic Shakespeare already may be among us, raising the same issue.

This issue is at the core of one of the subjects on the World Intellectual Property Organization's agenda for this conference--namely, "what the intellectual property status of the outputs of such [artificial intelligence] systems (such as computer-produced 'creations') should be?"[3] Others in this symposium are far more qualified technically than I to address the degree to which recent advances in artifical intelligence permit the creation of so-called "computer-generated" works with minimal human involvement. I will address whether we have entered a brave new world of copyright in which these works should be deprived of copyright protection because it may be difficult to identify a human author.

That WIPO has included this subject on the conference's agenda might suggest that it is a new one. But I submit that it is not. Rather, it has been with us since people began considering the impact of computers on copyright; moreover, upon analysis, the subject is not as troublesome as some suggest. We are not yet in a world of copyright without human authors, and there is no reason to believe that we are en route to that world. Nor is there reason to believe that if that world came about, it would create insoluble problems.

Answering the question of whether computer-generated works are copyrightable also requires answering other questions. The most obvious one is: if a computer-generated work is to be protected by copyright, who is the author? In the brief discussion that follows, I only can allude to some of the possible solutions to that question.[4] An equally important but (for me) relatively easy question is: if using a computer to create a work does not render it ineligible for copyright, what requirements must it meet to be eligible? The simple answer is that it must meet the requirements that any other copyrighted work must meet. For example, it must be an original work and usually--as in the United States[5]--a work fixed in a tangible medium of expression. If the computer independently creates "Ulysses," the computer-generated version will enjoy as much protection as James Joyce's version.[6]

Although the day on which a computer creates a "Ulysses" may not be approaching, artificial intelligence already has enabled computers to contribute to the creation of what most of us certainly would recognize as works of art, music, and literature. In some cases, the contribution of the computer clearly is limited to assistance. For example, choreographers can use computers to notate or preserve their dances and to experiment with new positions and movements;[7] but in those cases, the computer is simply being used to assist the human author in visualizing the author's creation.

In other cases, the computer's contribution may be more conspicuous and the human element of authorship less so. Computers have been programmed to create original works of art that genuinely please the eye.[8] A computer can take a soprano melody and, applying rules of musical composition, produce bass, alto, and tenor parts in conventional musical notation that harmonize with the melody, with a result that "sounds like good classical music."[9] Computers even have been known to produce poetry.[10] In those situations, the work of art, music or literature seemingly is created by the computer itself, although in reality the computer simply is following instructions that human programmers and users have given it.

Less esoterically, computerized "expert systems"[11] can produce solutions to a variety of problems by using an "inference engine" to process and manipulate a database ("knowledge base") prepared from the expertise of specialists in a given field. Whether the solutions are coyrightable works will depend on how the expression of those solutions measures up to other traditional copyright criteria.[12] If the expert system merely offers a simple diagnosis consisting of a few words (for example, an automotive diagnostic system which states the conclusion: "CLOGGED FUEL FILTER is confirmed"), its output will not rise to the level of a copyrightable work.[13]

Natural languages are another area in which the advances in artificial intelligence have implications for authorship. Computer programs are being developed to translate from one language to another.[14] When a work is translated by a human being, the translation itself can be a copyrightable derivative work.[15] Should the fact that a computer did the translating in accordance with human-produced rules disqualify the translation from copyright protection? Other natural language programs in development will prepare abstracts of articles or understand and "speak" in human languages, permitting lay users to "converse" with the computer. Who is the author of the computer-generated abstract? Is the computer the author of its half of the conversation with its human user?

My conclusion is that when such a work is otherwise entitled to copyright, it should not be disqualified because of the computer's contribution to the expression in the work. Copyright should not depend on the existence of a literal human author of the computer's output; the purposes of copyright law are served when an otherwise-copyrightable work is granted protection whether or not it was created with the intermediation of a computer. However, consideration of the question probably is nothing more than an interesting side excursion on the journey to our actual destination: it appears premature to consider the status now of a work of expression that is truly the product of a computer's "mind." Indeed, perhaps we still can ask in good faith whether that type of creation ever will come. Today's "computer-generated" works still have identifiable human authors, although it may not always be a simple matter to select that author. Therefore, for the foreseeable future, it suffices that the human element in the creation of otherwise protectible "computer-generated" works will sustain their copyrightability.

II. The History of the Issue in the United States

I can attest to the fact that authorship of computer-generated works is not a new issue, because I have been involved in at least two chapters of this question's relatively lengthy history. Twenty-six years ago, when the notion of artificial intelligence was primarily a gleam in the eye of some computer scientists, the United States Register of Copyrights pondered whether computer-generated works were entitled to copyright protection. Register Abraham Kaminstein reported that in 1965 the Copyright Office had received copyright applications for an abstract drawing and for compilations that were "at least partly the 'work' of computers," and that the Copyright Office previously had received an application for registration of a "musical composition created by computer."[16] the Register did not reveal whether the applications had been accepted[17] but he did offer this formula for resolving the issue:

> The crucial question appears to be whether the "work" is basically one of human authorship, with the computer merely being an assisting instrument, or whether the traditional elements of authorship in the work (literary, artistic or musical expression or elements of selection, arrangement, etc.) were actually conceived and executed not by a man but by a machine.[18]

Ten years later, in what were still--in the modern history of computers--virtually prehistoric times, the United States Congress established the National Commission on New Technological Uses of Copyrighted Works (CONTU) to address the copyright problems raised by computers that were not dealt with directly in the then-pending Copyright Act of 1976.[19] I was one of the commissioners appointed by President Ford. Among the issues that CONTU examined was the question of computer-assisted creation of copyrighted works,[20] a question whose genesis we traced back to Register Kaminstein's 1965 report.[21]

As I will relate in greater detail below, CONTU concluded unanimously that the artificial intelligence Register Kaminstein had envisioned ten years earlier had not yet been developed and did not appear to be on the horizon. The computer, like a camera or a typewriter,[22] was simply a tool to assist a human being in creating a work. As a result, there was no need to confront Register Kaminstein's crucial question of human authorship vel non, because there always would be a human author employing the computer and its program to do his or her bidding.[23] CONTU also concluded that in most cases, the author of a work created with the assistance of a computer would be the user of the computer.

The user was raised again, albeit obliquely, 1983, when Congress was considering the Semiconductor Chip Protection Act of 1984.[24] In the hearings before the Subcommittee on Patents, Copyrights and Trademarks of the Senate Committee on the Judiciary, Senator Mathias asked the representatives of the Semiconductor Industry Association whether it was possible to program a robot to design a semiconductor chip and if so, whether that chip would be the creative work of a human mind that was entitled to constitutional protection under the Copyright Clause of the United States Constitution.[25] Senator Mathias predicted that this question would raise a lot of interesting questions and that "the lawyers will thrive."[26]

I happened to be the next witness at the hearing, and I observed that "behind every robot there is a good person."[27] Recalling CONTU's conclusions, I observed that generally human intermediation precedes the conduct of the robot, or any other computer-based activity. Finally, I suggested that the words "writings" and "authors" in the Copyright Clause of the United States Constitution can be construed to embrace mask works and their creators, even if their formal production is aided by machine. I concluded that the fact that a work is partially produced mechanically with rays of light and chemicals should not deter us from recognizing its copyrightability, noting that we long had recognized the copyrightability of works produced by cameras, tape machines, and computer graphics.[28]

The issue was revisited yet again in 1986, when a report of the United States Congress's Office of Technology Assessment suggested that CONTU's approach may have been too simplistic. The OTA staff viewed computer programs as more than "inert tools of creation," and asked whether in some cases the programmed computer is at least a "co-creator" of a creative work--thus leading to possible difficulty in determining who is the author (for example, the user or the programmer, or both).[29] In a separate discussion of computer databases, the report noted a possible question as to whether "information" that a computer writes or compiles is a work of authorship, and expressed concern that copyright for machine-produced works might be a departure form copyright's traditional role as an incentive for authors.[30]

III. The History of the Issue Outside the United States

This summary of some of the history of this issue in the United States has counterparts in a number of other legal systems that have addressed the issue. For example, the Commission of the European Communities already has given the issue preliminary consideration. It said in its Green Paper that computer-generated computer programs should be entitled to copyright protection and that since the programmed computer is essentially a tool, those who use the programmed computer should be entitled to the copyright in its output.[31]

As described below, the United Kingdom has addressed the issue even more definitively, enacting a copyright statute that makes it irrelevant whether a computer-generated work can trace its origin to a human author. It provides instead that the author is deemed to be the person who undertakes the arrangements necessary for the creation of that work.[32] The approach proposed by the EC, in contrast, presumed that there is ultimately a human author.

The World Intellectual Property Organization, in the discussions of its Model Copyright Law, also is considering the status of "computer-produced works." At its third session last July, the WIPO Committee of Experts considered a MCL provision drafted by the International Bureau of WIPO that would define a "computer-produced work" as:

> "a work that is produced by means of computers, where the identificaion of the various creative contributions and the authors thereof is impossible [because of the number or the indirect nature of those contributions] [because the contributions of the authors are merged in the totality of the work]".[33]

The proposed WIPO draft provides that the owner of moral rights and the original owner of economic rights in a computer-produced work be either the person or entity "by whom or by which the arrangements necessary for the creation of the work are undertaken," or the person or entity "at the initiative and under the responsibility of whom or of which the work is created and disclosed."[34] The authors of the draft were of the view, however, that to qualify for protection under the Berne Convention[35] these works must trace their origin to a human author.[36] When the Committee of Experts considered this proposed provision, it concluded that further study was needed.[37]

IV. The Expandable and Flexible Nature of Copyright

Does this experience suggest that there is a consensus that computer-generated works should be entitled to copyright protection? The fact that government agencies with responsibility for examining the impact of computers on copyright law generally have concluded that protection should exist does not necessarily answer the question. Indeed, the fact that the issue has been addressed and examined on more than one occasion may suggest a degree of uncertainty.

Why does uncertainty exist? The idea of recognizing copyright in a work "created" by a machine may be unpalatable to those of us who are used to thinking of the traditional human author at his desk or artist at her easel. But as I already have suggested, even a so-called computer-generated work is not devoid of human authorship. People are involved in the creation of these works.

Our discomfort with the notion of computer-"authored" works (even if we cannot articulate a principled reason for the discomfort) is in keeping with a recurring phenomenon in the development of copyright law. In every age, a new technology has appeared about which people have expressed fear and concern, claiming that it defies the boundaries of the existing legal system. With respect to copyright, these claims were made about photographs, motion pictures, sound recordings, radio, television, and other telecommunications.

In each case, the copyright system has managed over time to incorporate the new medium of expression into the existing framework. Most recently, the role of the upstart new technology has been assumed by computers.[38] For a while, the computers-and-copyright battlefield was centered on the copyrightability of computer programs as literary works. That contest now has been largely fought and resolved in favor of copyrightability. It may be that the next battle will be over copyrightability of computer-generated works.

In his landmark article on copyright, one of my distinguished predecessors as Harvard Law School's professor of copyright, Professor Zechariah Chafee, Jr., characterized copyright as the "Cinderella of the law." He observed:

> "Her rich older sisters, Franchises and Patents, long crowded her into the chimney-corner. Suddenly the fairy godmother, Invention, endowed her with mechanical and electrical devices as magical as the pumpkin coach and the mice footmen. Now she whirls through the mad mazes of a glamorous ball".[39]

When Professor Chafee wrote these words in 1945, the magic of the technology with which copyright law contended was primitive compared with the wizardry of today's computers. It therefore is easy to understand the difficulty some have in accepting that copyright law applies to such mystical creations. Have the developments of the past 15 years outstripped the bounds of our current copyright law?

I think not. As I have suggested, copyright law has a built-in flexibility that enables it to adapt to the damands of new technologies. When the United States revised its Copyright Act in 1976, we recognized that they would continue to evolve; as a result, the new Act was drafted with flexibility. CONTU was created at the same time to make certain that those developments would not make the new ACT obsolete; we on the Commission concluded that surprisingly few amendments were required in order to take account of the implications of computers. By and large, we felt the existing framework was adequate. In fact, in passing the Copyright Act of 1976, Congress acknowledged that the "history of copyright law has been one of gradual expansion in the types of works accorded protection," including new forms of creative expression (such as electronic music and computer programs) that never had existed before, but had been made possible ny new scientific discoveries and technological developments.[40]

Professor Paul Goldstein, who shares this portion of the program with me, may have best described the flexibility of copyright law and its ability to adapt to technological change in his testimony before Congress about the Office of Technology Assessment report.[41] Discussing the challenges generally posed by present and emerging technologies, Professor Goldstein wisely stated:

> "I believe that the challenges presented differ little--certainly not in kind, and only slightly in degree--from the challenges that such technologies as radio, television, motion pictures, semiconductor chips--and, indeed, the printing press--have posed in the past. We <u>have</u> been there before. Thus, I believe that history and established principle offer the surest guides to Congress in resolving issues at the intersection between copyright and the new technologies."[42]

Although he was not specificially addressing copyrightability of computer-generated works, I believe Professor Goldstein's observations apply equally to these works. Computer-generated works may comprise one of the latest phenomena of new technology that copyright law must address, but indeed we have been there before, time and again. In the United States, for example, copyright has evolved over the last 200 years from a system offering protection only to maps, charts, and books[43] to a more flexible approach that protects all original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device.[44]

In light of this tradition of an expanding copyright law, why is there debate over whether computer-generated works are entitled to copyright protection? Do we hesitate out of nothing more than a visceral feeling that technology has thrown us a curve we are not prepared for? Or is there some principled basis for hesitating before welcoming computer-generated works into the family of copyrighted works?

The answer is the latter; and the principle that requires us to hesitate can be found between the lines of the section in the Copyright Act of 1976 quoted above, describing the subject matter of copyright. Is a computer-generated work a "work of authorship"? Register Kaminstein recognized this as an issue in 1965 when he stated the view that the crucial issue was whether the "work" is basically one of human authorship, with the computer merely being an assisting instrument, or whether a machine, and not a person, actually conceived and executed the traditional elements of authorship in the work.[45]

V. Common Law and Civil Law Approaches

The initial question,[46] then, is whether a human author should be required in order for a work to qualify for copyright protection. In the United States, our copyright law derives from the British tradition, with its pragmatic approach to copyright and authorship. In contrast, the Continental tradition may place greater emphasis on a human author,[47] whose absence may be less harmonious with the civil law copyright system.

Those of us in the Anglo-American tradition recognize that there is some difference between our copyright laws and those of the Continental or civil law tradition regarding the role in copyright of the individual's personality. This is evident in the doctrine of moral rights, which originated in civil law systems, although it has been transplanted to some degree to the British and United States systems.[48] Moral rights are justified as protecting the author's right of personality,[49] which at first blush appears to be a difficult notion to apply to computer-generated works.

On the other hand, the civil law tradition commonly considers moral rights to be separate from economic rights,[50] which we consider to be the core of copyright law. Economic rights, such as the exclusive right to exploit the copyrighted work, do not appear to be based on a notion that the work is a reflection of the author's personality. So, if a work were created by a nonhuman author, at the very least it would not appear that protecting it would be incompatible with the economic aspects of civil law copyright.

The Berne Convention is neutral on the subject. Article 1 states that the member countries constitute a Union "for the protection of the rights of authors in their literary and artistic works."[51] However, the Convention does not define "author." The Guide to the Berne Convention tells us that this is because "national laws diverge widely, some recognizing only natural persons as authors, while others treat certain legal entities as copyright owners."[52] United States law, for example, has recognized since the beginning of this century that an employer can be an author of works for hire.[53] The civil law system can reach a similar result, at least with respect to corporate ownership (if not authorship) of copyrights.[54] Japanese copyright law, despite having historical roots in the civil law system, provides that an author may be a natural or legal person[55] and recognizes a work for hire doctine in many ways similar to that of the United States.[56]

The Anglo-American tradition need not be troubled by these concerns. With its pragmatic emphasis on the economic aspects of copyright, it is not preoccupied with metaphysical notions of the relationship of the copyrighted work to the "personality" of its creator. Indeed, the current copyright law in the United Kingdom contains a separate provision for computer-generated works. "Computer-generated" is defined as meaning "that the work is generated by computer in circumstances such that there is no human author of the work."[57] With respect to such works, the UK law prescribes:

> In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken.[58]

Apparently unfettered by any fundamental requirement that a work must have human author, the British copyright law simply takes the pragmatic, commonsense approach that the absence of a human author is no impediment to copyrightability. If a work is created by a computer rather than by a person, the law simply will search for the human being who is responsible for the computer's creation of the work, and confer the copyright upon that person.

VI. "Writings" and "Authors" in United States Law

A. The Copyright Clause of the United States Constitution

In the United States, we may not have as clean a slate upon which to write as do the British. Our written Constitution permits the federal government to exercise only those powers delegated to it by the States under the federal Constitution. Thus, the national government's power to regulate copyrights is prescribed by the terms of the Copyright Clause.[59]

In determining whether a nonhuman author suffices under the Copyright Clause, it is fruitful to look to the express purpose of the clause: "to promote the progress of science and useful arts ..." This means that the exclusive rights of copyright are granted as an incentive to authors to produce works of authorship for the benefit of society.[60] The clause realistically recognizes that if an author knew that once the book was written, the picture painted, or the song composed, anyone could come along and reproduce it or otherwise exploit it, the incentive to produce the work would be diminished. There must be some assurance of a reward for creative labors (or at least that whatever reward there is will be the author's).

The computer, of course, needs no incentive to produce its output. A computer will do so for one reason: it is programmed to produce and told when to do so. What purpose would be served by granting the benefits of copyright to works if they are created by computers? Can we justify granting protection to them?

The computer may not need incentive to produce, but it will not produce anything unless human beings enable it to do so. Somebody has to write the program for the computer. And somebody has to instruct the computer to run a particular program that will produce the particular work. Somebody--or a combination of somebodies--ordinarily will have to give the computer precise instructions as to what to do. A computer may appear to create a work of art or music with no more prompting than the push of a button,[61] but usually there will be some prior human interaction with the computer, an interaction that shares in the "authorship" of the final product.

An incentive is just as appropriate for the human "collaborators" with the computer as for the starving artist or the impecunious writer. The computer may not go on strike if its output does not receive copyright protection, but it will not create anything unless there are people who care enough to prepare its programming and operate it. The difference between caring and not caring may depend on whether they can expect to enjoy the benefits of copyright.[62]

The Copyright Clause's purpose of promoting the progress of science and useful arts is hardly less served if achieved through computers, or human beings in collaboration with computers, rather than humans alone. The goal is "progress." It is not necessary to discard historic notions of the artist or author in her garret or study; recognizing computer creation simply acknowledges that valued works may be created under vastly different circumstances.

Thus, the policies underlying the constitutional basis for copyright are served by recognizing that works that may be created by computers are subject to copyright. The only other possible constitutional basis for disqualifying these works from protection would be the language of the Copyright Clause,[63] which permits Congress to confer copyrights upon "authors." Does the reference to "authors" literally require a human being? Or, does the United States Constitution permit the copyright law to protect a work that does not have human author? I incline to think that the answer to the first question is no, and, to the second, yes.

We know that our statute, which is enacted pursuant to the Copyright Clause and presumably consistent with it, explicitly permits, and our courts implicitly accept, legal entities to be considered as authors (as, for example, does the Japanese law[64]). True, our courts never have had to address that issue directly, but I do not believe that the few seemingly relevant phrasings in the pre-computer case law should lead us to conclude that an author must be a human being.

B. <u>The Case Law</u>

In one early decision, <u>The Trade-Mark Cases</u>,[65] in which our Supreme Court characterized the Copyright Clause as protecting "writings of authors," is acknowledged that the term "writings" may be construed liberally, but noted that only writings that are "original, and are founded in the creative powers of the mind" may be protected. "The writings which are to be protected are the fruits of intellectual labor."

The Court's views were expressed, however, in the context of the protectibility of a trademark. It said that our first federal trademark statute could find no constitutional basis in the Copyright Clause or the Patent Clause since a trademark "does not depend upon novelty, invention, discovery, or any work of the brain. It requires no fancy or imagination, no genius, no laborious thought."[66]

From this early description of the type of works protected by the Copyright Clause, one might conclude that a "computer-generated work" would not qualify, since it may be difficult to consider it a product of the "creative powers of the mind" or a "fruit[] of intellectual labor." To use other phrases of the Court, it might seem not to be "work of the brain," and to require no fancy, no imagination, no genius, no laborious thought. Yet, these are precisely the attributes that artificial intelligence consists of—as one definition puts it—"[t]he capability of a device to perform functions that are normally associated with human intelligence, such as reasoning, learning, and self-improvement."[67]

It would be dangerous to rely on the dicta in The Trade-Mark Cases to support any requirement that a copyrighted work be authored by a human being. When the Supreme Court decided it in 1879, the justices were not considering whether a computer, or any other non-human, could be an author. The question would not have occurred to the Court. The issue was not presented,[68] and no more preoccupied the Court than it did the drafters of our Constitution.

The passages from The Trade-Mark Cases were recently quoted by the Supreme Court in Feist Publications, Inc. v. Rural Telephone Service Co., Inc.[69] In holding that the white pages of a telephone directory did not have sufficient originality for copyright, the Court relied on The Trade-Mark Cases' discussion of the constitutional scope of "writings," concluding that originality requires "independent creation plus a modicum of creativity."[70] It found this requirement in The Trade-Mark Cases' references to protected "writings" as those that are "founded in the creative powers of the mind" and are "the fruits of intellectual labor."[71]

However, the century-old statements should not be read as implying the need for a human author. Feist was not concerned with who or what the author was; the issue was whether the compilation was the "writing" of an "author." In fact, a compilation created by means of artificial intelligence presumable passes muster under Feist so long as there has been a "modicum of creativity" in its selection, coordination or arrangement.[72] Feist does not appear to offer any opinion whether the requisite minimal creativity must be that of a human being or may be that of a computer programmed directly or indirectly by a human being.

The Court came a little closer to the issue five years after The Trade-Mark Cases, in the celebrated Oscar Wilde photograph case. Burrow-Giles Lithographic Co. v. Sarony.[73] The issue was whether a photographer could claim a copyright in his picture of the playwright. The defendant had reproduced lithographs of it, and contended that a photograph is not a writing of an author. The Court disagreed, and held that the photograph was a copyrightable work.

The Court did not hold that all photographs are copyrightable. It took pains to emphasize Mr. Sarony's authorship of the particular picture. It noted that "entirely from his own mental conception,"[74] he posed Wilde in front of the camera, selected and arranged the costume, draperies, and other accessories in the photo, and situated Wilde so as to present graceful outlines, all with the appropriate use of light and shade, to suggest and evoke the desired expression. It said that he had produced the photograph from that disposition, arrangement, representation, "made entirely by the plaintiff."[75]

Burrow-Giles is of interest because it implicitly raises the issue whether a camera can be an author. The lithographic company had argued that a photograph is a mere mechanical reproduction of a natural phenomenon and does not "embody the intellectual conception of its author, in which there is novelty, invention, [and] originality"[76] as required by the Copyright Clause. The Court acknowledged that "[t]his may be true in regard to the ordinary production of a photograph,"[77] but expressly declined to reach the broader question raised by the defendant. The reason it did not have to reach the question was that, as already noted, Mr. Sarony had shown that he had done more than simply snap the shutter of his own camera. He has proved "the existence of those facts of originality, of intellectual production, of thought, and conception."[78]

The Court therefore avoided the issue of whether a work created by a machine--in that instance a camera--was entitled to copyright. It resorted to a dictionary and expressed the view that an "author" in the sense envisioned in the Copyright Clause is "he to whom anything owes its origin; originator; maker; one who completes a work of science or literature."[79] It noted that at the time the United States Constitution was drafted, its framers understood the nature of copyright as the "exclusive right of a man to the production of his own genius or intellect."[80] The Court concluded that "the Constitution is broad enough to cover an act authorizing copyright of photographs, so far as they are representatives of original intellectual conceptions of the author."[81]

It is tempting to read this language as stating a requirement that a copyrighted work be created by a human author. However, it is doubtful that in these isolated passages the Court was even considering that question, let alone answering it. The Court was dealing with the technology of its time, and hardly with the question whether a machine might be capable of "intellectual conceptions."[82] There are limits to literal reading. For example, when it referred to an author as "he" to whom the work owes its origin and referred to the work as the product of "his" intellect, the Court was reflecting the mores of its time, and not dealing with whether women would qualify as authors.[83] Simply put, the Court was no more excluding machines from eligibility for authorship than it was excluding women. In other words, there may be less than meets the eye in this language.

Although Burrow-Giles reserved judgment on the issue, the lesson of its progeny is that, as a practical matter, virtually all photographs are eligible for copyright protection. Indeed, Learned Hand, one of our most distinguished jurists, particularly with regard to copyright, took the position in Jewelers' Circular Pub. Co. v. Keystone Pub. Co. that all photographs are protected by copyright, because "no photograph, however simple, can be unaffected by the personal influence of the author, and no two will be absolutely alike."[84] Moreover, he noted, the copyright statute protected photographs "without regard to the degree of 'personality' which enters into them."[85]

Judge Hand's use of the term "personality" may appear reminiscent of the civil law tradition of copyright; but his rejection of that notion as a copyright requirement reveals that, at least in Anglo-American jurisprudence, little deference is given to any requirement that a copyrighted work reflect an author's personality. The photographs in Jewelers' Circular were hardly works of art; rather, they simply were straightforward photographs of jewelers' trademarks for use in a directory. Indeed, had the photographer attempted to inject his personality into the photographs, he might well have been fired.

Other cases in our copyright jurisprudence could read, at first blush, as implying that an author must be a human being. For example, Goldstein v. California, a more recent Supreme Court case, interpreted the "writing" required by our Copyright Clause to include "any physical rendering of the fruits of creative intellectual or aesthetic labor."[86] But the Court in Goldstein did not deal with whether the "labor" of a computer system also would qualify to produce a "writing," nor with whether a "writing" could embrace anything beyond what the Court said it "included." On the definitional question, it concluded only that the "writings" under the Copyright Clause were not limited merely to script or printed material, but broad enough to include the Court's language quoted above and, specifically, sound recordings of artistic performances.[87]

Indeed, the Court reached its conclusion on this point only after emphasizing that "writings" and "authors" are terms that

> have not been construed in their narrow literal sense but, rather, with the reach necessary to reflect the broad scope of constitutional principles.[88]

Goldstein thus exemplifies the pragmatic core of our copyright system. Sound recordings (or "phonograms") generally are protected under copyright in Anglo-American jurisprudence.[89] To understand United States law, it is not necessary to debate the characterization sometimes given these works in some other jurisdictions as lacking the "personality of an author."[90]

Alfred Bell & Co. Ltd. v. Catalda Fine Arts, Inc.,[91] is another case touching upon considerations possibly relevant to whether a human contribution is required. In that case, in which the defendant challenged the originality of the plaintiff's mezzotint reproductions of old master paintings, a distinguished court of appeals said that to satisfy the constitutional requirement of originality, all that is needed is for an "author" to contribute something more than a "merely trivial" variation, something recognizably "his own."[92] But the language, although perhaps implying the contribution would be a human one, did not deal with that issue.

One judicial opinion did express a view on the subject, but its precendential value, if any, is quite limited. In Apple Computer, Inc. v. Franklin Computer Corp,[93] the district court concluded that to enjoy copyright protection, a computer program must be in a "language" that human being would understand. The court observed:

> If the concept of "language" means anything, it means an ability to create human interaction. It is the fixed expression of this that the copyright law protects, and only this. To go beyond the bounds of this protection would be ultimately to provide copyright protection to the programs created by a computer to run other computers. With that, we step into the world of Gulliver where horses are 'human' because they speak a language that sounds remarkably like the one humans use.[94]

The district court judge was speaking in a case in which there was no issue of a work created by a computer, and his comments were dicta. Moreover, his decision was reversed by the court of appeals.[95] Hence, if and when the issue of copyrightability of computer-generated works does come before a court in the United States, it may be unwise to rely on one judge's reasoning, just as it is unwise to rely on any of the other isolated bits and pieces of language in judicial opinions of the past century, none of which really addressed the issue.

It is far from clear that our courts would conclude that there is a requirement of human authorship. The reference in the Copyright Clause to "authors" does not tell us that these need be flesh and blood. Indeed, it tells us little more than that "authors" are the ones responsible for creating the "writings" that Congress is to protect. Two centuries ago, that meant only maps, charts, and books.[96] Today, of course, it means the broadened spectrum of works utterly unknown at that time, such as computer programs, computer databases, sound recordings, motion pictures (and we can go on and on). There is no reason why "authors" cannot undergo a comparable transformation.[97]

It does not appear that our copyright law is prevented constitutionally from protecting works created by computers. The underlying copyright policies do not require any such limitation; indeed they suggest the contrary.

VII. CONTU's Analysis of the Issue

Having said all that, I suggest it probably is unnecessary to reach a final resolution of the question whether copyright is available for works that are truly computer-generated, although we should continue to be mindful that it may well be answered affirmatively if it becomes necessary to answer it. The fact is that, as CONTU concluded thirteen years ago, it appears that even "computer-generated" works have human authors. It may not be a simple matter to determine who the author is, since there often will be competing candidates. But the possible difficulty of identifying the specific author is no reason to jettison the copyright protection that a computer-generated work needs and deserves as a work of authorship. CONTU's unanimous conclusion is still apt:

> Works created by the use of computers should be afforded copyright protection if they are original works of authorship within the Act of 1976. Consequently no amendment is needed.[98]

Cynics will suggest this passage seems to leave open whether these works in fact are "original works of authorship"--perhaps a typical example of question-begging by a governmental commission. Not so. Yes, the phrasing could have been more felicitous, but we explained our conclusion and made it clear that we found it difficult to conceive of any "computer-generated works" that would not qualify as original works of authorship.[99]

We were assisted in this conclusion by expert advice that the development of "artificial intelligence," which could create works copyrightable but for the lack of a human author, "has not yet come to pass." Thus, we noted that "indeed, it has been suggested to this Commission that such a development is too speculative to consider at this time."[100]

We concluded that a computer or computer program did not itself contribute authorship to a work produced through its use. Rather,

> [t]he computer, like a camera or typewriter, is an inert instrument, capable of functioning only when activated either directly or indirectly by a human. When so activated it is capable of doing only what it is directed to do in the way it is directed to perform.[101]

Although computers are complex and powerful, "it is a human power they extend." Thus, the computer "affects the copyright status of a resultant work no more than the employment of a still or motion-picture camera, a tape recorder, or a typewriter."[102]

CONTU made the assumption that copyright requires at least minimal human creative effort. It reviewed some of the cases discussed above, which addressed the originality requirement common to most systems of copyright law. From these cases, we concluded:

> Thus, it may be seen that although the quantum of originality needed to support a claim of authorship in a work is small, it must nevertheless be present. If a work created through application of computer technology meets this minimal test of originality, it is copyrightable.[103]

It appeared to us then, as it appears now, that there would be a human ingredient in any computer-generated work. This was not, as one now might suppose, merely a reflection of the relatively primitive times in which we were working. We were aware that artificial intelligence programs existed that, for example, would "select a series of notes and arrange them into a musical composition, employing various tonal qualities and rhythmic patterns."[104] We knew that programs existed through which computers could manipulate statistical information to produce an analysis of it that bears little similarity to the original form or arrangement of the work being analyzed.[105] In other words, we came to our conclusion notwithstanding our realization that computers could do things that, at first glance, did not appear to involve human creativity.

We concluded that computers are tools--very powerful ones, to be sure. They are tools used by human beings, and, and without human intermediation they do nothing. Progress has been marked by a continuing stream of creative developments that help people to take short-cuts in the process of production, including the creation of products of the mind. Although computers may be an unprecendented step forward in this progress, it seemed clear to us "that the copyright problems with respect to the authorship of new works produced with the assistance of a computer are not unlike those posed by the creation of more traditional works."[106]

In a sense, what CONTU did is reminiscent of what the Supreme Court did in the Oscar Wilde case[107] almost a century earlier. Just as that Court did not have to determine in Burrow-Giles whether a photograph that is not the product of a human being's "mental conception" is copyrightable, CONTU did not have to determine whether a computer-generated work with no human involvement is copyrightable.[108] In both instances, the conclusion was that there is a human author who is involved in creating the work; it is not just a mechanical operation. As a result, it is relatively easy to find that the work is entitled to copyright protection.

VIII. Has Anything Changed Since CONTU?

It is not 13 years after the CONTU Final Report. Is the development of a truly independent, creative artificial intelligence still speculative to consider? Or, have we progressed far enough to make wise policy decisions?

Others--not I--have the technical expertise to speak to the question of how much progress has been made. However, I am reminded of a comment by a specialist in one branch of artifical intelligence--neural networks--that the level of "intelligence" with which we have been able to endow such networks approaches (at best) that of the garden slug. To the extent that is the case with artificial intelligence in general, we will have to wait until computers have crawled further up the evolutionary scale before they can lay claim to the sole authorship of the wondrous copyrightable works they produce.

If, on the other hand, we have reached the point at which we can speak of computers or computer programs as the sole "authors" of original works, it hardly would be the first time that non-humans have been recognized as authors. As noted earlier, both the American and Japanese copyright laws provide that an employer that is a legal entity--a corporation or other institutional form--can be considered an author. Moreover, even if civil law jurisdictions may be less familiar with the concept of a corporate "author," it is not uncommon for some of these systems to recognize corporate ownership of copyrights.[109]

In any event, it should require little, if any, adjustment in most copyright systems to attribute the authorship to the human being who uses the computer, even if the computer is responsible for most or all of the work's creation. We already have seen how photographs may fit into this category.[110] Many photographs, such as Mr. Sarony's of Oscar Wilde, are works of art imbued with the photographer's personality, but it equally is clear that many photographs are simply mechanical productions of natural occurrences, as Burrow-Giles had argued. The photographs of the jewelry trademarks in Jewelers' Circular[111] are one example. Andy Warhols's film "Empire," in which he pointed a fixed motion picture camera lens at the empire State Building for eight hours, may be another.[112]

When Abraham Zapruder happened to point his home movie camera at President John F. Kennedy's motorcade in Dallas, he had no idea that it would record the assassination of the President. Yet a court upheld his copyright in the now-famous film,[113] even though the court acknowledged that it was "sheer happenstance"[114] that Zapruder was taking pictures at the scene and that he had started the camera "not knowing the horror it would record."[115] The defendants argued that the films were "simply records of what took place, without any 'elements' personal to Zapruder."[116] The court, however, followed Learned Hand's reasoning[117] that any photograph reflects the personality of the author, and overruled the objection to the film's copyrightability.

Another example is the hypothetical question I have put to the law students in my copyright classes: if a person walking down the street with a camera around the neck trips, accidentally releasing the shutter, is the resulting photograph a "work of authorship," subject to copyright? Although I still have the same unease about the answer as when I first posed the question thirty years ago, the authorities I have discussed suggest that the answer is yes, the photograph is copyrightable.

In such a case, it might make just as much sense to speak of the camera as the author. All the photographer did was start the camera running (and perhaps even that was an inadvertent or fortuitous act). This may be even more true with today's computerized autofocus cameras, which do everything but press the shutter release. Increasingly, it is difficult to consider many photographs as works that in any way are affected by the personality of the author.

Is the "photographer" who does nothing more than release the shutter any different from the computer operator who directs the computer to run an artificial intelligence program? I suspect that in most cases, the computer operator is engaged in a far more intricate intellectual process than the photographer. If the photograph is a copyrightable work of authorship, there appears to be no good reason why the computer-generated work should be treated differently.

IX. Who is the Author?

The real question may not be whether there is a human author, but rather who that author is. A number of intriguing possibilities come to mind. CONTU concluded that the "obvious answer is that the author is one who employs the computer."[118] We acknowledged that the simplicity of this approach may obscure some problems, but not problems appreciably different from determining authorship of other types of works.

For example, CONTU noted that often a number of persons are involved in using a computer to prepare a work such as a complex statistical table. Obviously, some are involved to a greater degree, and some in different ways, than others. Under United States law, when the collaborators prepare the work for a common employer, allocating authorship usually is not a problem, because the work would be considered a work for hire, with the employer as the author[119]--just as, for example, under the Japanese copyright law.[120] Other civil law systems can reach a similar result by considering the work a collective work owned by the person or legal entity "under whose name it is disclosed."[121] CONTU noted that if the collaborators do not have a common employer, conventional copyright law principles of joint authorship would apply. Thus, in the latter instance, the work may be considered a joint work, and the authors its co-owners.[122]

But as artificial intelligence programs become increasingly sophisticated and the task of the user increasingly simple and ministerial, it may be necessary for the search for a human author of a computer-generated work to go beyond the person or persons who "employ" the computer. Further refinement may be required: is the person who "employs" the computer simply the one who enters the command to run the program that creates the work? Is it the person who instructs the operator to enter the command? The person who decides to use the computer to create a particular work? The person who owns the computer and/or the computer program used to create the work? It may well be some combination of all of the above.

It is beyond the scope of this paper to come up with an all-purpose answer or formula to determine who the author of any computer-generated work may be. Numerous attempts have been made to resolve the issue.[123] Among the suggested resolutions are: apportionment between the user (problem-specifier) and the owner of the artificial intelligence software copyright;[124] apportionment between the person who created the program and the person who compiled the data;[125] the author of the underlying computer program;[126] the user of the program;[127] the computer itself;[128] the person who made the capital investment in the computer and its program;[129] or a fictional, nonexistent human author, with ownership of the copyright apportioned among various claimants.[130]

Commentators thus have differed as to who should be considered the author of a computer-generated work. They seem to agree that it should be some human being or legal entity,[131] even though identifying that author may not always be a simple matter. In theory, it would be convenient to have a bright-line rule that gives fully consistent and predictable answers, but we may have to settle for a more ad hoc, case-by-case approach that examines the specific facts of a given work and the various contributions to it.[132] Or, as is frequently the case when more than one person claims authorship of a work, the practical answer likely may be found in the agreement (or the intentions) of the parties.[133]

We need not be dismayed that the identity of the author will not always be clear. We have been in that ambiguous situation before as well. For example, when the United States Congress amended the 1909 Copyright Act to include sound recordings within the subject matter of copyright,[134] it observed that a sound recording involves authorship of both the artists whose performance is captured and the record producers who set up the recording session, captured and electronically processed the sounds, and compiled and edited them to make the final sound recording. Rather than make the difficult decision as to authorship, Congress recognized that the identity of the author varies in different situations. It therefore expressly decided that the statute should not fix the authorship or ownership of a sound recording, but should "leav[e] these matters to the employment relationship and bargaining among the interests involved."[135] Thus, there is significant precedent to conclude that recognizing authorship in a computer-generated work to support a copyright does not require that a general rule be formulated for identifying the author.

X. Conclusion

It seems that we have not reached a point at which artificial intelligence can be considered so "intelligent"--whatever that means--that an artificial author is creating a copyrightable work. We therefore must conclude for now, and the foreseeable future, that the copyrightability of otherwise protectible computer-generated works can be sustained by the human element in their creation, even though there may be difficulty in allocating authorship among various claimants. There probably is no more need today to resolve Register Kaminstein's question about computers as "authors" than there was when CONTU discussed the matter. However, if the day arrives when a computer truly is the sole author of an original work of art or music or literature (whether a novel or a computer program), copyright law will be sufficiently broad and adaptable to recognize these productions as among its protected works.

NOTES

1 However, the issue of animal authorship actually may arise in other contexts. For all we know, "Little Dropout," the monkey in the Dick Tracy comic strip who painted by splattering on a canvas, may have produced works of art rivalling those of Jackson Pollock. There is no reason why "Little Dropout" could not have his counterpart in the real world. The American TV entertainer Johnny Carson, among others, delights in featuring dogs who bark to music, and the dogs often appear to be composing their own music.

2 In this paper, I refer to the production of works by "computers," by "computer programs" or by "computer systems." The significance in each instance is the absence (to a greater or lesser degree) of the appearance of direct human authorship of the output; in this context, I use the terms interchangeably.

Similarly, there are overlaps between terms such as "computer-produced," "computer-generated," "computer-created," and "computer-assisted" works. Although the copyright principles we apply to them are the same, the terms connote varying degrees of human involvement. An example at the very low end of the spectrum is the word processing program (including a utility to check spelling) that was used in the preparation of this paper, but which does not make the computer the author. At the other end of the spectrum presumably would be a computer that would by itself conduct the research, engage in the reflection and write the paper.

3 World Intellectual Property Organization, _WIPO Worldwide Symposium on The Intellectual Property Aspects of Artificial Intelligence; General Information and Provisional Program 3_ (Doc No. SAI/INF/1, November 12, 1990).

4 _See_ the discussion in text at notes 118-135, _infra_.

5 _See_ Section 102(a) of the Copyright Act of 1976, 17 U.S.C. 102(a).

6 For this paper, I shall assume that the computer-generated works I question meet all other requirements for copyright, with the possible exception of, arguably, human authorship. As with works created in more traditional ways, some computer-generated works will satisfy those requirements; some will not.

7 _See_ J. Dunning, _Dance by the Light of the Tube_, The New York Times Magazine 26 (February 10, 1991).

8 _See, e.g.,_ K. Sofer, _Art? Or Not Art?_, Datamation 118 (Oct. 1981); P. Garrison, _Glued to the Set_, Harvard Magazine 26 (January 1989).

9 _See_ T. Maugh II, _Researcher Uses Bits, Bytes and Bach for Program of Note,_ Los Angeles Times, August 10, 1988, part 1, p.30, col. 1.

10 M. Borroff, _Computer as Poet_, Yale Alumni Magazine 22 (January 1971). Borroff used a "random stanza" program to produce several stanzas of poetry. She provided the computer with 19 separate vocabularies of 50 words each. Each of the vocabularies served a particular function (for example, verbs in the imperative mode, or nouns in the vocative case). A series of "frames" was

constructed, consisting of grammatical function-words such as articles and prepositions, including blank slots into which words selected at random from designated vocabularies were to be inserted. Thus, Borroff contributed much of the authorship to the poetry generated by the computer, with "authorship" contributed by the computer in the form of random selection from a limited vocabulary preselected by Borroff.

Presumably, there would be little difficulty in concluding that Borroff was the actual author of the poetry created by the computer. However, one can imagine that a descendant of Borroff's program might be capable independently of generating Shakespeare's sonnets. In that case, perhaps the programmed computer itself would be the author.

11 A relatively succinct definition of "expert system" may be found in A. Freedman, The Computer Glossary 229 (1991):

> A[n] artificial intelligence application that uses a knowledge base of human expertise to aid in solving problems. The degree of problem solving is based on the quality of the data and rules obtained from the human expert...
>
> The expert system derives its answers by running the knowledge base through an inference engine, a software program that interacts with the user and processes the results from the rules and data in the knowledge base.
>
> Examples of expert systems are medical diagnosis, equipment repair, investment analysis, financial, estate and insurance planning, route scheduling for vehicles, contract bidding, counseling for self-service customers, production control and training.

12 See note 6, supra.

13 See J. Pepper, An Expert System for Automotive Diagnosis, in R. Kurzweil, The Age of Intelligent Machines 330-335 (1990). The Service Bay Diagnostic System described by Pepper asks a series of questions about the car's condition (for example: "Does the car crank?"; "Is there spark at the ignition coil?"). The programmed computer evaluates the answers by referring to the knowledge base, which contains rules devised by human diagnosticians (for example: "If car doesn't crank, remove NO FUEL from possible causes for NO START"). Applying these rules, the expert system reaches a diagnosis (for example: "CLOGGED FUEL FILTER is confirmed").

See also 1 M. & D. Nimmer, Nimmer on Copyright 2.01 [B] (1990) ("Nimmer"), discussing the quantum of originality required by copyright. While that quantum is low, see the discussion in text at notes 67-94, infra, nevertheless it rules out protection for fragmentary words or phrases and forms of expression dictated solely by functional considerations. Nimmer at 2-14 through 2-15. "Courts are disinclined to permit copyright to attach to short word sequences or to find plagiarism in the copying of such sequences." B. Kaplan, An Unhurried View of Copyright 46 (1967). Nor are the individual diagnoses of such expert systems likely to pass the test that "a short phrase may command copyright protection if it exhibits sufficient creativity." Nimmer at 2-16.

14 See R. Kurzweil, supra n. 15, at 307-312.

15 See, e.g., the definition of "derivative work" in Sec. 101 of the United States Copyright Act of 1976, 17 U.S.C. 101; Olympia Press v. Lancer Books, Inc., 267 F. Supp. 920 (S.D.N.Y. 1967).

16 Copyright Office, Sixty-Eighth Annual Report of the Register of Copyrights 5 (1965).

17 Id. See also Copyright Office, Sixty-Seventh Annual Report of the Register of Copyrights 4 (1964), in which the Register reported on Armstrong Cork Co. v. Kaminstein, No. 119-64 (D.D.C., filed January 16), a mandamus action seeking to overturn his decision to refuse registration of a design for flooring when the patterns were produced haphazardly as a result of vinyl chips falling at random through a hopper. The Register had concluded that the "design" did not constitute the "writing of an author." Id. at 5.

18 Copyright Office, SUPRA n. 18, at 5.

19 Pub. L. 93-573, title II: National Commission on New Technological Uses of Copyrighted Works (1975).

20 CONTU was assigned the task of studying and compiling data on the creation of new works by the application or intervention of computers, as well as recommending any changes in copyright law or procedure necessary to preserve public access to such works, and to recognize the rights of copyright owners. See CONTU Final Report at 43-44. This assignment came in the wake of questions about computer-generated works during the debates on the new copyright law. Part of our mandate was to determine whether there should be any amendments to the law relating to those works.

21 CONTU Final Report 44-46 (1979).

22 It is not necessary to claim for CONTU either prescience or omniscience in order for me to call attention to the sophisticated cameras and typewriters ("word processors") of today, which are truly complex computer systems of hardware, software, and "firmware" microchips.

23 CONTU Final Report at 43-46. See the discussion in text at notes 100-110, infra.

24 The Semiconductor Chip Protection Act ("SCPA"), 17 U.S.C. 901-914, established a quasi-copyright regime for the protection of semiconductor chip mask works. There was considerable debate whether mask works should be protected by a copyright approach (favored by the Senate) or by a sui generis approach (favored by the House of Representatives). The sui generis approach prevailed.

I do not think the sui generis approach was correct for mask works, and I would question whether it would be appropriate for the broad spectrum of computer-generated works, which may include traditional works of art, music, literature, audiovisual works, choreography, and others—all of which otherwise qualify for traditional copyright protection.

25 U.S. Const., Art. I, Sec. 8, Cl. 8:

The Congress shall have Power...To promote the Progress of Science and useful Arts, by securing for limited times to Authors...the exclusive Right to their respective Writings...

26 The Semiconductor Chip Protection Act of 1983: Hearings on S. 1201 Before the Subcomm. on Patents, Copyrights and Trademarks of the Senate Committee on the Judiciary, 98th Cong., 1st Sess. 86 (1983). The question was posed following the Statement of a Panel Consisting of Thomas Dunlap, Jr., Corporate Counsel and Secretary, Intel Corp., Accompanied by Richard Stern, Copyright Counsel; and Christopher K. Layton Vice President of Operations, Intersil, Inc., General Electric Co., on behalf of the Semiconductor Industry Association, Accompanied by Stanley C. Corwin, patent Counsel.

27 Statement of Arthur R. Miller, Professor of Law, Harvard Law School, in id. at 88.

28 Id. at 88. I offer my conclusions and contributions of several years ago not with any belief that they are the last word on the subject, but simply to illustrate that this issue is not a new one, and has been discussed repeatedly over the years.

29 Office of Technology Assessment, Intellectual Property Rights in an Age of Electronics and Information 70-73 (1986).

30 Id. at 76. See, however, the discussion in text at notes 62-64, infra (regarding incentive), and the discussion in text at notes 64-69 and 111-119, infra (regarding authorship).

31 Commission of the European Communities, Green Paper on Copyright and the Challenge of Technology, Doc. No. COM (88) 172 Final, 5.6.25 & 5.6.26 (1988).

The subsequent Proposal for a Council Directive on the Legal Protection of Computer Programs, Doc. No. COM (88) 816 Final-SYN 183, Official Journal of the European Communities (89/C91/05) (12 April 1989), would have provided that for programs generated by the use of a computer program, the natural or legal person who causes the generation of subsequent programs shall be entitled to exercise all rights in the program, unless otherwise provided by contract. Id., Art. 2(5). However, the recent Common Position Adopted by the Council on 13 December 1990 with a View to the Adoption of a Directive on the Legal Protection of Computer Programs (Doc. 10652/1/90) is silent on the issue. See id., Art. 2.

32 See the text at notes 57-58, infra, and the paragraph thereafter.

33 Preparatory Document, WIPO Draft Model Law on Copyright, No. CE/MPC/III/2, 126 (March 30, 1990). The two bracketed clauses represent "two alternatives to indicate the reason why it is impossible to identify the various creative contributions." Id., 127.

34 Id., 258-259. The first alternative was based on the UK provision. The second alternative was based on a provision proposed by WIPO governing "collective works." The drafters solicited comments as to which alternative would be more appropriate.

35 Berne Convention for the Protection of Literary and Artistic Works (Paris Act 1971).

35 Preparatory Document, supra n. 35, 122. It is difficult to reconcile this view, however, with the contrary interpretation of Berne set forth in the authoritative WIPO Guide -o the Berne Convention and exemplified in the copyright laws of a numbe: of Berne members--for example, Japan, the United States, and the United Kingdom. See notes 55-60, infra, and the accompanying text.

37 Report Adopted by the Committee, No. CE/MPC/III/3 72-76, 134 (July 13, 1990).

38 See, e.g., CONTU Final Report at 27-37, in which Commissioner John Hersey dissented from CONTU's recommendation regarding copyrightability of computer programs and argued that the Copyright Act should be amended to make it explicit that copyright protection does not extend to a computer program in the form in which it is capable of being used to control computer operations.

39 Z. Chafee, Jr., Reflections on the Law of Copyright: I, 45 Colum. L. Rev. 503 (1945).

40 H.R. Rep. No. 1476, 94th Cong., 2d Sess. 51 (1976).

41 Note 29, supra.

42 Statement of Paul Goldstein before the Subcommittee on Patents, Copyrights and Trademarks, Committee on the Judiciary, U.S. Senate and the Subcommittee on Courts, Civil Liberties, and the Administration of Justice, Committee on the Judiciary, U.S. House of Representatives 1 (April 16, 1986).

43 Public Act of May 31, 1790, 1st cong., 2d Sess. 1.

44 Copyright Act of 1976 Sec. 102(a), 17 U.S.C. 102(a). For a brief synopsis of the historical expansion of the scope of copyright in the United States, see Goldstein v. California, 412 U.S. 546, 562 n. 17 (1973), in which the Supreme Court noted that "[a]s our technology has expanded the means available for creative activity and has provided economical means for reproducing manifestations of such activity, new areas of federal protection have been initiated," and the Court outlined the increasingly expansive revisions of the copyright statutes in the United States.

45 Note 16, supra.

46 Although the question is an important one, we may never reach it because, as discussed below, there is reason to believe that (now and in the foreseeable future) the authorship of a computer-generated work can be traced to one or more human beings.

47 Indeed, this is suggested by the fact that the French term for copyright is "droit d'auteur," and other languages use similar terms.

48 See, e.g., Copyright, Designs and Patents Act 1988, Ch. IV (United Kingdom); Visual Artists Rights Act of 1990, Pub. L. No. 101-650, 104 Stat. 5089, 5128, 17 U.S.C. 106A (effective June 1, 1991) United States).

49 See P. Geller, International Copyright: An Introduction: in International Copyright Law and Practice Int-28 (P. Geller ed. 1990); R. Plaisant, France, in International Copyright Law and Practice at Fra-12 (P. Geller ed. 1990).

50 See id. However, this may not be the case with respect to German law. See P. Geller, International Copyright: An Introduction, in International Copyright Law and Practice Int-28-29 (P. Geller ed. 1990). See also Article 2(2) of the German Act dealing with Copyright and Related Rights (as amended up to June 24, 1985), in Copyright Laws and Treaties of the World ("'Works' within the meaning of this Act include only personal intellectual creations" (Emphasis added)).

51 Paris Act, 1971, Art. 1.

52 World Intellectual Property Organization, Guide to the Berne Convention 11 (1978).

53 The 1909 Copyright Act provided that "the word 'author' shall include an employer in the case of works made for hire." An Act to Amend and Consolidate the Acts Respecting Copyright, March 4, 1909, 60th Cong., 2d Sess., 62. The current law, the 1976 Act, provides:

> "In the case of a work made for hire, the employer or other person for whom the work was prepared is considered the author for purposes of this title..."

Section 201 (b) of the Copyright Act of 1976, 17 U.S.C. Sec. 201(b).

54 See, e.g., Article 9, 3rd par. and Article 13 of the French Law of March 11, 1957 on Literary and Artistic Property; R. Plaisant, France, in International Copyright Law and Practice at Fra-42 (P. Geller ed. 1990).

55 Chosakuken-ho (Copyright Act), Law No. 48 of 1970, Articles 2(1)(ii), 2(6). See T. Doi, Japan, in S. Stewart & H. Sandison, International Copyright and Neighbouring Rights 783 (1989).

56 Chosakuken-ho (Copyright Act), Law No. 48 of 1970, as revised by Law No. 62 of 1986, Articles 15(1) and 15(2). See T. Doi, supra note 57, id.

57 Copyright, Designs and Patents Act 1988, Sec. 178.

58 Id., 9(3).

59 U.S. Const., Art. I, Sec. 8, Cl. 8, set forth in note 27, supra.

60 The Supreme Court has recently observed that "[t]he primary objective of copyright is not to reward the labor of authors, but '[t]o promote the Progress of Science and useful Arts.'" Feist Publications, Inc. v. Rural Telephone Service Co., Inc., 499 U.S. ____, 113 L.Ed.2d 358, 59 U.S.L.W. 4251, 18 U.S.P.Q. 2d 1275, 1279, 18 Med.L.Rptr. 1889, No. 89-1909, slip op. at 8 (March 27, 1991).

61 And even in those cases, it is fair to conclude that a human ingredient of the creative process has been largely carried forward from an earlier stage of the process--for example, from the creation of the original program or the database used by the program in a given instance.

62 The human collaborators may receive these benefits directly or indirectly--for example, as individual entrepreneurs, or in compensation from legal entities for whom they have produced works for hire or commissioned works. These entities in turn are provided incentive for their investment in research and development and in the dissemination necessary to bring these works to society.

63 Note 25, supra.

64 See notes 55-56, supra.

65 100 U.S. 82 (1879).

66 Id. Uncopyrightable trademarks apparently must be contrasted, however, with copyrightable photographs that can be made of such trademarks by a machine that we call a camera--an ironic contrast presented by Judge Learned Hand's decision in Jewelers' Circular Pub. Co. v. Keystone Pub. Co., 274 F. 932 (S.D.N.Y. 1921), aff'd on other grounds, 281 F. 83 (2d Cir.), cert. denied, 259 U.S. 581 (1922), upholding the copyrightability of photographs of trademarks. See the text at notes 86-87, infra. The Court of Appeals in Jewelers' Circula did not address the copyrightability of photographic illustrations, affirming instead on the ground that the work was protectible as a directory or compilation as a result of "industrious collection" of the source material by the author. That ground for protection was expressly disapproved by the Supreme Court in Feist, 18 U.S.P.Q.2d at 1280-81, slip op. at 11-12.

67 D. Spencer, The Illustrated Computer Dictionary 8 (1980).

68 The issue was whether legislation regulating trademarks could be based on the Copyright Clause.

69 Note 61, supra.

70 18 U.S.P.Q.2d at 1278, slip op. at 5.

71 Id.

72 Feist acknowledges that the creativity required is minimal:

> "To be sure, the requisite level of creativity is extremely low; even a slight amount will suffice. The vast majority of works make the grade quite easily, as they possess some creative spark, no matter how crude, humble or obvious it might be. [Citing 1 M. Nimmer & D. Nimmer, Copyright 1.08 [C] [1] (1990)]" 18 U.S.P.Q. at 1278, slip op. at 4.

73 111 U.S. 53 (1884).

74 Id. at 54-55.

75 Id. at 55.

76 Id. at 58-59.

77 Id. at 59.

78 Id. at 60.

79 Id. at 57.

80 Id. at 58.

81 Id. Feist, supra n. 62, discussed these passages from Burrow-Giles in its discussion of originality, characterizing the case as "emphasiz[ing] the creative component of originality." 18 U.S.P.Q.2d at 1278, slip op. at 5. However, for the reasons already discussed (see the discussion in the text accompanying notes 71-74, (supra), Feist appears not to present any view that a human author is a sine qua non of copyright.

82 Referring to the language in several older copyright decisions, including Burrow-Giles, Professor Goldstein correctly observes that "these dicta, expressed at a time when no court contemplated the potential of the modern computer, are hardly authority for a requirement of direct human authorship." I P. Goldstein, Copyright 2.2.2 at 73 (1989).

83 Likewise, as to the Court's references to the nature of copyright as the exclusive right "of a man" to the production of "his" own genius or intellect.

84 274 F. 932, 934 (S.D.N.Y. 1921), aff'd, 281 F. 83 (2d Cir.), cert. denied, 259 U.S. 581 (1922).

85 Id. Nimmer questions Judge Hand's reasoning to the extent it suggests that the 1909 Copyright Act exempted photographs from the requirement of originality. M. & D. Nimmer, supra n. 15, 2.08[E] [1]. However, the treatise accepts Judge Hand's view that all photographs are affected by the personality of the author and concludes that almost any photograph "may claim copyright merely by virtue of the photographer's personal choice of subject matter, angle of photograph, lighting, and determination of the precise time when the photograph is to be taken." Id.

86 412 U.S. 546, 561 (1973). The Court relied at this point on Burrow-Giles and The Trade-Mark Cases.

That interpretation could be read, as CONTU read it, to imply a limitation to human labor only. CONTU Final Report at 45. However, CONTU did not cite Goldstein as holding that our system requires human authorship. Goldstein was referred to in a discussion of the requirement of originality, in which CONTU noted that only a modicum of effort was required.

87 412 U.S. at 562.

88 Id. at 561.

89 See, e.g., 17 U.S.C. Sec. 102(a)(7).

90 The United States Congress has confirmed the American view that these are protectible "writings of an author":

> The copyrightable elements in a sound recording will usually, though not always, involve "authorship" both on the part of the performers whose performance is captured and on the part of the record producer responsible for setting up the recording session, capturing and electronically processing the sounds and compiling and editing them to make the final sound recording.

H.R. Rep. No. 1476, 94th Cong., 2d Sess. 56; see also H.R. Rep. No. 487, 92nd Cong., 1st Sess. 2, 5 (1971); S. Rep. No. 72, 92nd Cong., 1st Sess. 4 (1971); Goldstein v. California, supra, 412 U.S. at 561-562.

91 191 F.2d 99, 102-103 (2d Cir. 1951).

92 In fact, the court concluded that even if the author copying a pre-existing work unintentionally came up with a distinguishable variation (for example, due to the fright engendered by a clap of thunder), that variation would be sufficient to support a copyright so long as the "author" (the quotation marks were used by the court) adopts it as his own. Note that if true the variation would not have been intentional, indeed not even volitional.

93 545 F. Supp. 812 (E.D. Pa. 1982), rev'd 714 F.2d 1240 (3d Cir. 1983), cert. dismissed, 464 U.S. 1033 (1984).

94 Id. at 825.

95 714 F.2d 1240 (3d Cir. 1983), cert. dismissed, 464 U.S. 1033 (1984). The court of appeals did not express an opinion on the copyrightability of computer-generated works.

96 Public Act of May 31, 1790, Sec 1.

97 The dictionary definition of "author" is sufficiently flexible to include nonhumans. See, e.g., Webster's Ninth New Collegiate Dictionary 117 (1987) ["the writer of a literary work"; "one that originates or gives existence: SOURCE"]. Under the United States copyright law--and many others as well--"literary works" include computer programs and computer databases. See H.R. Rep. 1476, 94th Cong., 2d Sess. 54.

See also the discussion of Burrow-Giles, in the text at note 81, supra, in which the Supreme Court resorted to the dictionary to aid it in determining what is an "author."

98 CONTU Final Report at 1.

99 See id. at 43-46.

100 Id. at 44.

101 Id.

102 Id. at 45.

103 Id. at 45. We continued our conclusion, in words that (as discussed above) may have exaggerated the need for human authorship:

> The eligibility of any work for protection by copyright depends not upon the device or devices used in its creation, but rather upon the presence of at least minimal human creative effort at the time the work is produced.

If this was an overstatement of the law, it was excusable at the time as CONTU's gloss on Alfred Bell & Co., Ltd. v. Catalda Fine Arts, Inc., 191 F.2d 99 (2d Cir. 1951), which the Commission discussed in a passage just preceding the quoted language. See CONTU Final Report at 45.

104 Id. at 44.

105 Id.

106 Id. at 45.

107 Note 73, supra.

108 CONTU implicitly did leave open the possibility that at some point in the future, artificial intelligence could independently create works that have no human author, although that development was "too speculative to consider at this time." CONTU Final Report at 44. We recognized that "the dynamics of computer science promise changes in the creation and use of authors' writings that cannot be predicted with any certainty" and recommended that the situation be monitored. Id. at 46.

109 See notes 54-56, supra, and the accompanying text.

110 Protection of photographs regardless of artistic merit is not unknown in civil law copyright systems as well. See, e.g., R. Plaisant, France, in International Copyright Law and Practice at Fra-21 (P. Geller ed. 1990).

111 Note 84, supra.

112 Warhol, of course, did select the film, the location and time of the filming, and possibly the lens, filter, exposure, and other elements.

113 Time Inc. v. Bernard Geis Associates, 293 F. Supp. 130 (S.D.N.Y. 1968).

114 Id. at 131.

115 Id. at 133.

116 Id. at 141.

117 Id. at 141-143. See Jewelers' Circular, supra n. 86.

118 CONTU Final Report at 45.

119 Se note 53, supra, and the accompanying text.

120 See note 56, supra, and the accompanying text.

121 See note 54, supra, and the accompanying text.

122 17 U.S.C. Sec. 101, 201(a).

123 See, e.g., T. Butler, Can a Computer be an Author? Copyright Aspects of Artificial Intelligence, 4 COMM/ENT 707 (1982); P. Samuelson, Allocating Ownership Rights in Computer-Generated Works, 47 U. Pitt. L. Rev. 1185 (1986); E.H. Farr, Copyrightability of Computer-Created Works, 15 Rutgers Computer & Technology L.J. 63 (1989); B. Sookman, Computer Assisted Creation of Works Protected by Copyright, International Computer Law Adviser 8 (October 1989); S. Nycum & I. Fong, Artificial Intelligence and Certain Resulting Legal Issues, The Computer Lawyer 1 (Vol. 2, No. 5, May 1985); Report of the Committee to consider the Law on Copyright and Designs (Cmnd 6732) (1977) (the United Kingdom's "Whitford Report") 514-515; Milde, Can a Computer be "An Author" or an "Inventor"?, 51 J. Pat. Off. soc'y 338 (1969); D. Rosen, A Common Law for the Ages of Intellectual Property, 38 U. Miami L. Rev. 769, 804 (1984); S. Hewitt, Protection of Works Created by the Use of Computers, New L. J. 235 (March 11, 1983).

124 Nycum & Fong, supra n. 125. The authors note that "typically..., a user will direct the computer system by describing the problem to be solved or the style of composition desired." Id. at 6. They conclude that a court would apportion ownership between the owner of the AI software copyright and the user/problem-specifier, using traditional copyright analysis. However, they note that at the stage of technology current in 1985, the user normally would be the sole author even though he has taken advantage of computerized aids in the process.

125 The Whitford Report, supra n. 125 515. The Whitford Committee considered the computer to be a mere tool used in the creation of a work. Therefore, it is clear that the author of the output can be none other than the person, or persons, who devised the instructions and originated the data used to control and condition the computer to produce the particular result. In many cases it will be a matter of joint authorship. We realize this in itself can cause problems, but no more than in some other fields, and we are not convinced there is a need for special treatment.

126 E. Farr, supra n. 125,at 73-74, 79-80. Farr, discussing a hypothetical short story-writing program, reasons that

> although the programmer may not know exactly what the program will produce each time it is executed, nevertheless he is the one who supplies the vacabulary lists and the grammatical and syntactical rules that enable the program to produce its output.

Id. at 73. Farr concludes that it is the programmer's idea that is being expressed when the computer creates a work and that the programmer is "the only individual who contributes enough creative intellectual effort to satisfy the copyright requirement of authorship." Id. at 80.

See also D. Rosen supra n. 125, at 803-804. Rosen, however, may have been considering a situation in which the programmer is also the user. Rosen's example is computer artist Harold Cohen, who programs his computer to create works of art. See also K. Sofer, supra n. 10. Based on Burrow-Giles' statement that the issue is whether the work "owe[s] its origin" to the "author," 111 U.S. at 57-58, he concludes that

> [v]iewed in that light, ... artificial intelligence programmers are indeed the authors of their computers' works. Although the machines make decisions on their own, those decisions are made within confines established by the programmer/artist.

38 U. Miami L. Rev. at 804.

127 P. Samuelson, supra, n. 125, at 1200-1204, 1227-1228. Samuelson recognizes that in many cases the user will not have contributed sufficient authorship under traditional copyright analysis, but concludes that statutory and policy reasons favor allocating authorship to the user. The user is the "instrument of fixation for the work, that is, the person who most immediately caused the work to be brought into being." Id. at 1202. Moreover, the "originality" requirement of copyright has a threshold sufficiently low to admit users as authors. Id. Since the user presumably already has paid for the right to use the program (thereby giving reward to the programmer), the user should have the right to "us[e] the work for its intended purpose of creating new works." Id. at 1203. Further, the user often will have played a significant role in shaping the output of the computer. Id. at 1203-1204. Finally, giving copyright to the user is "the most practical solution and the one least likely to lead to litigation." Id. at 1227-1228.

128 Milde, supra, n. 125, at 393-395. Milde believes that when the computer exercises independent creation, it is at least a co-author (with the programmer and the person who provided it with data) of the work it creates. Id. He also concludes that when the output of a computer contains any "deviations" from its input, the output is "written" by the computer. Id. at 403.

129 S. Hewitt, supra n. 125, at 236-237. Hewitt reasons that the person who made a capital investment in these items "may reasonably expect some return." Id. at 236. Although he recognizes that the investor's claim to authorship is tenuous since he has not originated the idea, he finds the analogy to a work for hire attractive, with a "creative machine" taking the place typically occupied by an employee. Id. at 237.

130 T. Butler, supra n. 125, at 744-745. Butler considers attributing authorship to a fictitious human being to be "an alternative which avoids [the] pitfalls" of trying to find the real human being responsible for authorship of the computer-generated work. Courts called upon to determine authorship and ownership of the products of AI programs should

> presume the existence of a fictional human author and assign the appropriate fractions of the copyright rights to the owner of the AI software copyrights, the problem-specifier or the computer owner, either individually, jointly or in part.

131 Samuelson does suggest the possibility that no copyright protection should be given if no human author can be identified, but ultimately concludes that his "conflicts with the temper of the times" and that the user should be deemed to be the author. P. Samuelson, supra n. 125, at 1224-1228.

132 As Judge Learned Hand said about a similar question in a well-known copyright case, "[o]bviously, no principle can be stated as to when an imitator has gone beyond copying the 'idea,' and has borrowed its 'expression.' Decisions must therefore inevitably be ad hoc." Peter Pan Fabrics, Inc. v. Martin Weiner Corp., 274 F.2d 487, 489 (2d Cir. 1960).

133 Professor Goldstein notes that rights in computer-generated works can be allocated by relying on contractual arrangements between the copyright owner of the computer program and the user of the program. I.P. Goldstein, Copyright 2.2.2 at 73 (1989).

134 Sound Recording Amendment of 1971, Pub. L. 92-140 (1971).

135 H.R. Rep. 437, 92nd Cong., 1st Sess. 6 (1971). See also M. & D. Nimmer, supra n. 15, 2.10 [A] [3].

ARTIFICIAL INTELLIGENCE IN THE FILM INDUSTRY: AN OXYMORON?

by

Norman Alterman
Senior Vice President, Legal Affairs
Motion Picture Export Association of America (MPEAA)
New York, N.Y.
United States of America

You, the audience have sat and listened to a number of very erudite speakers on a variety of highly technical subjects. I come to you without footnotes or transparencies--armed only with some experience in my particular field of expertise, and, perhaps naively, totally lacking in the fear that normally generates caution when speaking to as learned a group as this.

I come to you as a proprietor, a user, and most of all, as a dreamer. I deal with applications and uses. There are no doubt a number of people listening to these words today who wonder at the connection between the film industry and a subject as complex as artificial intelligence. So many people in the world believe that the film industry, at least in America, is a stellar example of a lack of intelligence. Artificial, certainly. Artifice, perhaps. But, they believe, hardly intelligence.

The film industry is usually thought of as an industry based on dreams, puffery, artificial everything, imitations of life, but certainly not intelligence. Nonetheless, computers have a long and visible history in film making. A simple, memorable example. The graphic environment created for the Disney movie TRON, about the "people" in a computer game, was thought by some to be an example of advanced graphic techniques.

Lucasfilm showed how far that first film could be taken, through the STAR WARS trilogy. Others have further advanced the state of the art. In fact Lucasfilm itself has an entire division called Industrial Light and Magic devoted to meeting just such challenges.

The same computer graphic techniques produce the visual portions of the Disneyland attractions such as the Trip Through Space. Combined with the physical movement of the space in which the viewer is seated, one has the absolute illusion of a space trip. I, for one, am reminded of the first Planetarium show I saw as a child, long before the first large scale computer, ENIAC, was built. And, I am reminded how primitive was that marvel.

I. SOME PRESENT USAGES

A. WRITING

In this day and age it is hard to find people who write for a living who still use foolscap and fountain pen to produce lengthy documents such as scripts, or even treatments or outlines. Word processors are commonly found as the current recording tool of choice among those who turn in the written materials on which the films are based.

B. ANIMATION

As anyone knows who has explored computer graphics, animation is a natural outgrowth of the basic drawing programs. For some years now, much animation has been done in Japan by computer. To observe Japanese animation, it is only necessary to watch children's television programming on Saturday mornings in America. It is quite sophisticated.

C. COLORIZATION

That nasty word. That quarrelsome concept. Yet, without sophisticated computer programs, impossible economically if not physically. The quality of the colorization is directly proportional to the sophistication of the computer program.

D. GRAPHIC ENVIRONMENTS

My favorite synthetic environment is the "alley" through which Luke Skywalker had to pilot his craft in order to drop the bomb into the central nervous system of the Empire Death Star spaceship. I truly felt that I was there...in a real environment...in space.

An article in a recent issue of Smithsonian magazine described how far beyond STAR WARS are the current applications presently being created at the Industrial Light and Magic division of Lucasfilm Studios. There seems to be no limit to the possibilities.

This is a big step forward from the 360 degree "surround" picture developed by Walt Disney studios for the New York World's Fair, and now a permanent fixture at Disneyland.

II. POSSIBLE ARTIFICIAL INTELLIGENCE APPLICATIONS

A. WRITINGS

To the extent that it is possible to create a story or better yet a script by what is called artificial intelligence, then it would be possible to have applications in this area.

But, can anyone in the audience conceive of this happening? Are we back to the 500 monkeys seated at the 500 typewriters to see how long it would take to produce the collected works of William Shakespeare? I think so.

As has already been noted, a recent newspaper article reports a novel written in the style of Jacqueline Suzanne, by a collaboration between a good computer programmer and the program he wrote. While he certainly may not claim that Jacqueline Suzanne herself wrote it, he certainly can claim it to be an original creative work. But, should he get credit as the author of the book or as the author of the program or as the author of both the book and the program? Should he be subject to an infringement action for having replicated the "look and feel" of a Jacqueline Suzanne novel?

Could a story be converted into a script by an artificial intelligence system? Probably. But I firmly believe that the story must exist first, even if only in outline form as a concept. It seems inconceivable to me that anyone other than the person who manipulated the computer, even if only to turn it on, could be considered as the author.

B. PRODUCTION DATABASES

This may be the most productive current use of computers. A properly loaded database, regularly updated, can do an instantaneous match between the story description of the locale and a physical location where the film can be shot or where a second unit can do the background filming.

The budget, using the familiar spreadsheet principles with imported database information, can be instantly determined, and determined in the alternative for each variable. It can give relatively precise figures for each competing location, for each competing actor, director, producer, and could factor in past performance, assisting the producer in making intelligent decisions best calculated to produce a winning formula - one that wins both at the critical level and at the financial level.

How convenient to try out in advance the various possible locations available for the conversion of the story into a different point of view, or a different mood, and then decide based on the cost, the cost-benefit ratio of each possible variation.

C. DIRECTION

At present, given the limited state of my knowledge, the main use here would seem to be as an electronic notebook, or perhaps in advance planning for the actual shooting.

D. SET DESIGN/CONSTRUCTION

Here there is an obvious use of computer programs.

There are a number of programs currently available which deal with the layout of a specific site, as well as its decor. The program permits changing the layout and viewing it from all angles to determine desirability of a particular design. Such a program also permits decoration and redecoration of a given site.

Should the computer get screen credit? What about the author of the computer program? What of the set designer in all of this?

Query: can an audiovisual work be filmed in cyberspace?

E. COSTUME DESIGN AND MAKE-UP

Just as a site or set can be designed and decorated, so can the appearance of a person, even to the reproduction of the general shape and size of the person, the relationship of the person to other persons in the production, and such other uses as the human mind can conceive.

F. EDITING

Using a computer to keep track of all of the choices, the cuts, the sequences, the type of fades, dissolves, etc., desired by the editor would be easy with a computer, especially connected to multiple screens so that all footage under review at any given time were simultaneously visible. Then the computer could produce the sequence for the laboratory to follow in printing the first copy rather that have to laboriously perform the same operations manually. For material digitally recorded, the entire process can be done in a desktop computer.

III. SOME COMMON ELEMENTS

What are the common elements in the above uses? Is there any thread that runs through all of them?

Every one of the uses set forth involves the computer assisting a person in the doing of a task, even if that assistance is so extensive or so sophisticated to rise to the dignity of being called a computer-generated work. Who assigns the task? Who sets the parameters for the task? Who sets the goals for the action? The computer or the person? I believe it is the person. Who therefore should get the protection? The person not the computer. The computer does not program itself to perform the task. The computer does not select the goal or the desired result. The person does. In each case, does the computer initiate the action? Or terminate the action? Or decide whether such a termination is temporary or final? Of course not. The person does. The traditional creative spark must come from a person, even if his only decision is to start the computer or to run the program. This is true no matter what the level of sophistication of the program. The machine does not turn itself on unless it is programmed to do so upon the happening of a event or sequence of events pre-determined by a person.

Even if this approach doesn't resolve the riddle of the legal protection of artificial intelligence, there is a next level of intelligence. Let us say that the computer is a clone, or descendant of HAL 9000, the superior computer of the film "2001", and that it can program and re-program itself. Was it born that way? Is that ability genetic? Or is it contained in a chip? Or a program? The computer, no matter how "smart", does not mutate, nor does it evolve, in the biological sense, even if it could construct a clone of itself. It only exists as a series of "1's" and "0's", which do not themselves have a life cycle. Can the "1's" and "0's" rearrange themselves? Yes, if taught to do so.

Can an artificial intelligence system teach itself to do so? Yes, but only if the program to teach it to do so has been installed and activated.

Regardless of the level of sophistication of the operating program, that operating program had to be written, constructed, conceived. Could a machine such as a computer write, construct and conceive that operating program? Yes, but only if another program had been installed and activated which instructed the computer to write, construct and conceive the program. And so on, and so on, ad infinitum.

Ultimately, one must return to the fundamental operating program which had to be written by a human to produce the results.

I believe that, at the root of every sophisticated program which emulates the human mind, lies a human mind.

In addition, there are certain decisions that are not yet able to be made by computers. Perhaps one day they will. But I think it will be long time before decisions on whether, where, in what medium and when to release or publish a film will be made by a computer without human intervention. I say this despite the fact that one well-known film producer uses computers extensively to track distribution of his films and to assist him to make decisions about subsequent distribution and marketing. I sincerely doubt that computers will be substitutes for rather than aides to human judgment in the film industry, unless a human makes that decision for a particular film or group of films.

Would this create a new category of "superstar" computer as creator to demand a percentage of receipts as compensation?

Finally, should the author of the program be deemed a co-author or collaborator? Maybe, but let the parties themselves decide in the licensing agreement. Should the computer itself be considered as the author, or at least the producer, if it does all of the above and also arranges the financing? Perhaps, since in most cases, arranging the financing is the truly creative part!

IV. CASE STUDY

It is not necessary to project very far into the next century, and perhaps not even beyond this century, to envision a simulation which comes fairly close to an artificial intelligence system.

We know that computer dictionaries have been with us in ever more sophisticated forms for a number of years. We know that writing styles can be emulated. Combining the two, we now have the ability to translate a writing style into another language.

Let's add a complication. Suppose that the writing is a film script, and the translation is to be used for dubbing the film. Can the computer be programmed to match an acceptable translation of the text fairly closely to the lip movements of the characters on the screen? If we have an intelligent enough computer, or computer program, then of course the answer is "yes".

So, we now have a very sophisticated program that translates film scripts into other languages with minimum loss of lip synchronization.

Until my mythical year, the lines still had to be spoken by a real person. But, now we must have and therefore will create a program that will make the real dubbing performers unnecessary.

Envision a program that analyzes the voice of the original performer, and then reproduces that voice to speak the translated script. Ah, you say, it will sound like the artificial voices one gets when calling a computer programmed telephone number information service, speaking...one...word... at...a...time. Nonsense, say I. If I can envision a proper program that does all of the above, then I can also envision a program that also analyzes the speed of speech, tone, and emotional content of the original speaker, and puts it all back into the synthetic voice!

Artificial intelligence or no? I answer: Artificial, yes. Intelligent, no. A good trick? You bet! And hats off to the programmer or team of programmers who makes it happen.

In my mythical year, the marvelous voices of marvelous performers will still be mar-elous, or at least something that resembles their own voices, not the voice assigned for the film. Now, if we can do this for a translation, think what this technique would have done for the careers of all those great performers who disappeared because their voices couldn't make the transition from silent films to talkies. Where is John Gilbert? Rudolph Valentino? Vilma Banky? A host of others .... Would their films have current value if appropriate computer generated voices were added?

Is there an equally exciting use for this technology when it is developed? How about a satellite television transmission, in which a program goes up in one language and comes down in a language of the viewer's choice?

What would it take, a 30-second delay? Or is that too long? If it were a live transmission, how long would it take to make the necessary ephemeral recording needed to perform the computer translation and lip-synch? The rest of the 30-second delay I proposed originally? Would it be used? Would people pay for the privilege? How many other uses for such a program have each of you thought of while I described it? For the technical experts in the audience, how many of you are already working on some aspect of this concept? How long will it take to produce an economically viable and technologically and emotionally acceptable program? Five years? Ten years? Two years?

Would the program physically be located at the transponder? Would it be cheaper and easier to install it in someone's home? Would the black box then be language specific? Or could you translate for guests who spoke a different language? Would that require an added fee? Obviously this is not a frivolous question. Consider its use in hotel rooms ....

V. FUTURES

Now we can go back to examine some interesting possibilities of the use of "artificial intelligence" in the film industry.

If a person's voice can be synthetically produced, why not his likeness? How would you like to see a film starring John Wilkes Booth and Abe Lincoln? If today, a film about Julius Caesar were to be made, who better to play the part than Caesar himself? Since there is no other audiovisual record of such people, it is not necessary because there is no basis for comparison with any performer playing the role.

On a less fanciful plane, consider the plight of the unfortunate producer whose star suffers a fatal accident just before the final principal scenes are filmed. Now the film can be finished with a synthetic recreation. You are horrified, but is this any different from animation at an incredibly sophisticated level? Leaving aside the morality or even the moral rights, why not have the possibility especially if the synthetic recreation is so skillfully done as to be indistinguishable from the original?

If the original of a sculpture can be reproduced in this way through a good computer-assisted drawing program, why not a performance? If a single note in a sound recording can be replaced, why not a misspoken word or an unwanted expression in a film? What is the difference? Where is the line to be drawn?

And, once again, I submit to you that these are choices made in the end by people, not computers. The computers, the programs thought of as artificial intelligence, are all merely tools of the creative artist.

VI. THE MONKEYS AND THE TYPEWRITER

Even if the monkeys, working at random at their hypothetical typewriters, could write, even out of sequence, all of the plays of William Shakespeare, people will be unlikely to be willing to wait for the good stuff through the decades during which all of the gibberish writing is produced.

This kind of random selection, chance, is certainly enhanced by the ever-increasing ability to instantaneously examine all of the random chances.

But, the number of random chances is certainly reduced if the creative spark emanates from that greatest computer of all, the human mind, as inevitably all creation must.

So, what will be the future of computers and of "artificial intelligence" in the film business? I see an enormous use in the construction, care and "feeding" of databases of every type, purpose, description, shape, size, accessibility and complexity, that is in the assistance that can be rendered to the person responsible for the decision. I see a more limited use as a tool, albeit an extremely sophisticated tool, for the creation of dreams, the dreams of which films are made. And I see at the root of it all, an author, a real live biological human being.

INTELLECTUAL PROPERTY PROBLEMS RELATED TO ARTIFICIAL INTELLIGENCE SYSTEMS

by

Antonio Millé
Latin-American Institute for Advanced Technology
Computer Science and Law
Buenos Aires
Argentina

1. Introduction

The participants in the WIPO Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence, after having received for two days information on Artificial Intelligence from the most qualified experts, will understandably disdain this brief introduction considered necessary for the author who, only with his experience in intellectual property law and his limited practice as an occasional user of the PROLOG language, must plunge--with no previous warm-up--into his topic.

1.1. Object of study

I will refer to "artificial intelligence" (AI) as the information systems which carry out the capture and processing of data about reality, intended to solve physical and abstract problems by means of logical resources confined, until now, to the human mind. I agree that the contours of AI are not well defined. Therefore, functions covered by AI systems that can also be performed by information systems should not be defined as "Intelligent."[1]

To carry to an extreme the typical features of Informatics, AI "is a center of multiple convergencies: research on Symbolic Logic ... Cognitive Psychology ... Neuroscience ... Cybernetics,"[2] which originate additional inconveniences because they require treatment of juridical problems related to different fields of human activity, many of which have special juridical and regulatory problems. This paper does not intend to deal with all of them, but to concentrate on those typical of the intellectual property categories, closely related to AI systems since they are the result of the creative activity of human intelligence.

1.2. Importance and composition of the market

Undoubtedly, the issue in question has significant importance due to the number of creators devoted to the production of AI systems, the quantity and diversity of applications they offer to the community, and the economic range of the market for these creations. Such importance is growing, according to projections which indicate an increase of 12.878% for the U.S. domestic market of products and services related to AI in the decade 1983/1993.[3]

As to its composition, the world market estimated at $ 6,470 million[4] for 1990, was split in the following manner:

Expert Systems 34.16%
Vision Systems 17.08%
Computer Hardware Specific for Artificial Intelligence 15.52%
Natural Languages 12.42%
Artificial Intelligence Elements in Robotics 6.21%
Military Use 4.97%
Telecommunications, Air and Space 4.5%
Computer Assisted Design 2.79%
Voice Recognition 2.33%

This predominance of "expert" applications surely stimulates the greater attention legal commentators have paid to problems related to expert systems,[5] which, on the other hand, contribute with a richer material to intellectual property law, because they contain compiled information.

2. Assets to Protect

A preliminary analysis leads us to inquire about the different externalizations of AI that correspond to immaterial assets for which protection is required through intellectual property. I will propose the following categories:

(a) AI systems offered as a finished product, ready for its use for the purpose they are designed. Within this category, experience shows two classes:

(a)(i) Systems open to the user who can complete or modify them, by himself, after receiving them from the creator.

(a)(ii) Canned systems, commercialized in executable version, with unchangeable files.[6]

(b) AI administration systems, commercialized as "shell," intended for final users to create the knowledge bases which will allow to give them a functional objective.[7]

(c) Knowledge bases commercialized separately, for use under the command of inference engines of a third's authorship.[8]

(d) Immaterial assets produced, totally or partially, by AI. Among them, products that, while they were created by humans, received derived protection from:

Copyright;
Patent Law;
Trademark Law;
Trade Secrets and Know-How;
Industrial Models and Designs;
Semiconductor Chip Protection.

Covering my topic, I will review the problems related to the protection of categories "a" and "b," focusing the analysis on the main categories of intellectual property. This does not imply ruling out the possibility that the authors of AI systems resort additionally to self protection through trade secret or to limited (since it can not be exercised erga omnes) protection offered by bilateral agreements.

3. Possible Protection Resources

As in other cases[9] of creative complex manifestation, AI systems may be protectable through different categories of intellectual property. Such protections need not be exclusive, some of them pointing to the "essence" and others to the "form" of the intellectual activity result.

In the same way they are not exclusive, protection offered by the different categories of intellectual property law cannot be comprehensive: as to the legal protection we will distinguish within each AI system physical and logical components, some of whose elements will be protected by more than one type of intellectual property, others than only by one of them, and others by none.

At the same time, when functioning, AI systems can occasionally produce internal results (rules and self-learnt and/or self-elaborated knowledge) and/or external results ("works," "inventions," "trademarks"[10]) which can be protected, or not, by some of the intellectual property categories. Thanks to AI systems' capacity to gain experience and apply rules in an unexpected way for their author, most of these results have not been--nor could they be--projected by human beings who contributed to the production of processing or basic information resources of AI systems. In such cases, must we give those products of technological fortuity the same legal treatment we give to the natural products whose immateriality has no owner?[11]

As is visible, we are in front of a complex problem that can be perfectly solved through the existing legal institutions and the analysis methods already known, but that presents more inconveniences due to the multiple nature of the immaterial asset under consideration.

4. The "Sui Generis" Solution

The Law created protection resources for immaterial assets, as they started appearing in the market and as soon as social conflicts in relation to their property were detected. That happened within Gutenberg's invention, and more recently, when semiconductor chips appeared. Is it, then, convenient to think of another specific protection for AI?

The solution does not appeal to me at all. In a world that constantly incorporates new ways of producing immaterial assets and restlessly enlarges the possibilities of exploiting them, when more and more utilitarian articles offer to the public a mixture of material and immaterial components protected by different categories of intellectual property, the appeal for "ad hoc" protection in this case--or any other where it can be avoided--would only disrupt the unity of the intellectual property system.[12]

On the other hand, this solution has not been posed for the case, nor foreseen in the courrent tendencies of intellectual property.[13]

5. Protection Under Patent Law

Although the history of most recent judicial cases tend to question its absolute validity,[14] there is a worldwide legal principle that denies computer software its capability of being protected under invention patents. This will be a determining factor in the protection of AI systems, whose immaterial component is made up of organized data files (evidently unfit for patent) and software resources.

In that way, AI systems based entirely on immaterial resources, such as Expert Systems, will not see patent law present in their protection panoply. On the contrary, patent law will appear as a possibility to be taken into account for the protection of physico-logical systems of AI, where software controls certain processes, in the case of new products or new processes that meet the requirements of novelty and non-obviousness.

6. Protection as a Semiconductor Chip

Of course, there is a possibility that part of an AI system is made up of semiconductor chips, which have a specific protection--aimed at reserving the exclusive use of its physical structure--in the countries that have adopted such system. We must stress that this protection will complement, not eliminate, the one given by patent law to new inventions, where the semiconductor chips take part, and by copyright to the expression of ideas that the semiconductor chips support.

7. Protection by Copyright

7.1. Some applicable principles

Before expressing some ideas about when I consider the main legal resource for the protection of this kind of immaterial assets, I think it is timely to recall typical details of the discipline, which are useful to solve problems related to this kind of work.

- For works of high technical complexity, the issue of the legal preservation of "non literal elements" shows that the problem of protection for the work as a whole (something not always easy to achieve when what is really valuable belongs to the field of the "ideas") is definitely linked to the protection claimable for component parts of the work and for component elements of the work parts.[15]

- As to works of any kind and degree of technical complexity, we cannot determine in advance the protection of the category but only its fitness to benefit from protection if the general conditions required (originality, personal intellectual labor, etc.) are given in each concrete case.[16]

- Copyright protection differs from patent law in that the latter covers even the minimum details of the protected invention. Copyright, on the contrary, does not cover the work as a whole but only its expressive elements capable of being protected. But on the other hand, if such expressions are original and of non-compulsory use, they will benefit from legal defense and will not be required to have the novelty and inventive level needed for the patentable invention.

7.2. Essence and form

In AI, as in any other utilitarian creative manifestation, the expression serves the functional condition. That's why essence and form are very close and it is occasionally difficult to "disentangle"[17] the expression from the idea.

By applying the method of moving along the abstraction scale in order to distinguish a "pattern" of idea from its expressions: how many grades can we go through in the case of a clear and synthetically expressed "production rule"? Will there be any grade at all?

Another difficulty is typical of any scientific work, whose most valuable element lies in the ideological essence that cannot be protected.[18] Obviously it could not be expected that the knowledge or practical experience should receive, in its "essence," greater protection from the law when transferred to the memory of an AI system than when expressed directly by a human being.

Even if the problem of separation between the elements of "essence" and "form" (i.e. "expression" and "idea") is solved, and if confinement of the protection to the pure expression is accepted, the inconveniences do not end there: some AI systems, as neural networks,[19] have a form assigned to them by the system at the learning, that evolves as that learning continues. Therefore, there is neither a fixed representative form wanted and achieved by the author to express the idea, nor any certainty that products from two different authors will not adopt a similar form at some stage of its evolution.[20] Shall we then conclude that the neural networks--as to the expression through "weights" of the predicates--may turn out to be a variety of intellectual work generated by computer?

When these lines reach their readers, maybe other more acute colleagues will have presented their solution to this enigma in the course of the Symposium.

7.3. Data bases

In principle, the data gathered in data bases (whether facts or rules) will be doubly protected in the form of writings and compilations.

The protection of the text as a written work will be in danger many times due to its simplicity and briefness, a necessary consequence of its submission to syntax requirements. But if obligatorily summary[21] the expression used will not always be obligatory (there will be other ways of expressing the information: more or less redundant, more or less condensed) for which reason, in most cases it can be distinguished from the idea.

Even in the presence of texts belonging to the knowledge base where idea and expression are merged, the author must make up the data base through the careful selection of a certain sequence, structure and organization. Although this will not offer protection to every and each information considered individually, it will constitute a defense resource for the creation considered as a whole.

As to compilations, the knowledge bases will be affected by the typical problems of this type of work, among them the difficulty of protecting the exhaustive data bases. Such bases--because they shelter all the universe of the data in question--do not seem to reflect a selection criterion and therefore deprive themselves of one of the foundations for protecting compilations.

In addition, the knowledge bases make up collections of assertions that are often loaded without a certain order, as well as modified, added and removed freely, and then used according to the order the inference engine implements in each case to solve a concrete problem. Does this circumstance destroy the possibility of protecting a "sequence and organization" that does not derive from a creative intellectual human activity?

7.4. Inference engines

As to the inference engine and any other software resources, its use for AI does not add or remove anything: it is one more manifestation of the problem of the software's legal protection, already solved.

Maybe, what can create an interesting problem in the future are the "non-literal" elements of certain inference engines, especially those belonging to expert systems. The matter of the user's interfaces authoral reservation (in the process of being solved by copyright law) has in the case of AI systems a greater scope, not only related to the "look and feel," but also to strategies of questioning and help for the user that surpass the field of the mere expression.

7.5. Problems of Copyright Law related to AI

The reflection on the problems that will be originated by the application of copyright principles on AI systems,[22] allows to detect some areas of possible conflict.

7.5.1. Protection of languages

The creation of an AI system can include the production of a computer language (vocabulary and syntax) intended to express instructions or data. This matter will also present the problem of languages' protection, still not sufficiently debated by intellectual property law.

7.5.2. Authorship attribution

As in case of software, in almost any field of AI we find ourselves in front of works developed with the contribution (individualized and distinguishable or not) of a plurality of intellectual workers, whose part in the work shows different degrees of creativity related to the elements of "essence" and "form" in different ways and measure. Besides, this kind of work is carried out under the direction and responsibility of a corporation by people linked to it by a contractual or service relation.

The problems related to collaboration become more serious in works of this sort, due to the diversity of contributions. Experts on different arts and sciences, knowledge engineers and information systems people, will contribute to a common product. However, some contributions will always be of an "essence" nature, while others will consist of giving the "form" to that essence. Will this cause that any of the contributors is not considered a co-author? Or, on the contrary, must we think that because of the special nature of the creation, the co-authorship of the contributors as a whole should be considered, in principle, as existent?[23]

Note that in this sort of work there is a contributor who takes the "essence" from the interrogation to the expert and--after "formalizing" it at an intermediate stage--he usually hands it over to the information systems person, in order for the latter to produce its definitive "expression." Nevertheless, we recognize that the task of the knowledge engineers has such relevance that, without it, most of the complicated expert systems would not exist. It has been observed that the knowledge engineer many times obtains from the expert, information he already had and had not been conscious of.[24]

Will the joint use of an AI's "shell" and a knowledge base of a different authorship create a specific case of collaboration, giving way to the birth of a "composite work?"[25] We rather think that the "add-in's" make up independent works--that mutually complement themselves during the execution by the user--very different from the "collection compiled form pre-existing works."

7.5.3. Matters related to the "moral right"

As in the case of software, the main difficulty in relation to the moral rights of the author has to do with the right to oppose any unauthorized modification of the work--first, due to functional reasons, as these works require "maintenance" both corrective and updating; second, because--as we have already seen--the functioning of the work in itself and its collection of "experience" produces modifications in its literal expression.

We also detect a problem that can be classified under this item (though it is also related to personal rights and employment law) which has to do with creative freedom (something like the right to decide on disclosure of the work). This problem can emerge in the case of systems that "learn" from the operation by human beings, recording their practical work, operation sequence, options, etc., for which reason they can be applied to register the "know-how" of people not in accordance with--and occasionally ignoring--it.

7.5.4. Formalities

As has been already observed, some AI manifestations--particularly expert systems--show a changing "form," both as a result of the constant addition, modification and removal of elements in the knowledge base, and of acquisition of "experience" by the system itself. It was noted that such characteristic is stressed in the neural networks.

In this context, what will be the "expression" of the work for which a registration will be requested or whose copy will be deposited in the countries that still demand the fulfillment of formalities as condition for the validity and/or enforcement of the copyright? What protection will shelter the "stages" (since these are not truly "versions") that follow the application?[26]

7.5.5. The "use" of the work

In AI systems, we view the problem of "use" as typical of utilitarian works (as computer programs or data bases). This kind of work is exploited permanently and continously by the final user, in order to obtain concrete utilitarian results. Even though the use is related to reproduction (since the work could not be possibly utilized in the corresponding equipment if it had not been previously reproduced in copies) from the standpoint of the work's exploitation, the extension and characteristics of the use are determinant for the practical benefit the user will obtain and the economic reward the author can claim. In short: they are the kind of work from which editors, distributors and users regard more accurately their license contract than the copy itself.

It is then forseeable that this sort of work should present the same problems currently faced by computer programs, in particular ones related to the fulfillment of contracts and liability, to which I refer separately.

7.5.6. Contractual aspects

As a natural consequence of the fact that these works are "used," the terms of the copyright owner/user relationship grow in importance and we must expect that specific contractual structures are made common for this kind of licenses.

As "tailored systems" become marketed with a certain frequency of time, we will also incorporate in the repertoire of intellectual property law the corresponding standard formulas. Everything seems to show that this kind of contracts shall be strongly similar to the ones currently used for software development and license agreements.

As to the internal relations within the group of co-authors and between the principal who takes charge of the work and his dependent or hired employees, there is no doubt that the contracts will have a decisive relevance.[27] Such agreements will define not only the extent of the cession of intellectual property contributions and commitments of non-disclosure and confidentiality, but also other key aspects of the matter, such as those related to non-competition and non-employment.

7.5.7. Liability problems

The issue of liability for damages caused to third parties that used the work in accordance to its objective, is a matter traditionally outside the scope of copyright, to such an extent that traditional studies on the subject lack even the mere mention of the topic. However this is an issue the existence of multiple works with utilitarian purpose has made transcendent.[28] As far as copyright lacks specific rules, to the present the issue is solved on the basis of the general principles of the law.

The owners of rights in AI systems must be particularly careful and take precautions against this kind of hazards, since the product they market not only runs the risk of malfunctioning--typical of computer software--but also add the possibility of an error in the knowledge base (i.e. in the ideological essence of the work). On the other hand, expert systems takes charge of "non physical" works performed to the present by human intellect, whereas robots take up the material tasks until now performed by man. AI then faces a new problem: the legal consequences of decisions and actions taken under the control of a system.[29] So, it becomes indispensible to bear in mind the risk of possible liabilities of the system's creators and users to third parties.

In the same way as that of the author's liability for the consequences of his work's use, the "ethical" aspect (in the jargon of this specialty "moral" has a specific and different meaning) of copyright was, for long, considered an antique linked to censorship. But AI shows us the possibility of intellectual works that include the author's ethical (or amoral) concept, perhaps contrary to the user's or the social environment's one.[30] The matter deserves reflection and consideration, since it could bring about assumptions of criminal liability.

The hypothesis of errors capable of producing harmful consequences are numerous in AI systems and this produces liabilities. Among them:

- Errors in the "short term memory" data input (mistake made by the human operator or information badly grasped by the equipment designed for that function) or in the "long term memory" (when the human expressed a rule wrongly or ambiguously).

- Failures in the process (the inference engine did not work properly due to a faulty design or bad programming[31]).

- Mistake in the application of rules (a rule was missing or it produced failures in combination with others.)

- Errors in the output (ambiguity in the expression of the result, incorrect functioning of the equipment commanded by the system, etc.).

Apart from the "direct" errors of human authors, the capability of gaining experience, and the fact that AI systems can determine on their own actions to be followed, could also create liability:

- If the system was operated following the correct legal rules, but its result contradicted other legal principles based on a different logic. For instance, a system of personnel selection that, on the basis of rules supported only by appraisals of the job capability and vocation, ended up always rejecting candidates of a certain sex, religion, nationality or ethnic origin.

- If, as a consequence of its normal functioning, the system made up or used data bases containing personal data in a manner aginst the Personal Privacy Protection Law of certain countries.

- If the system carried out actions that would be unlawful in the assumed case they had a human agent. For instance, if an AI system uses information in a defamatory way or, is used as an EDI resource, it incurs actions objected to by the law of unfair competition or monopoly.

- Lastly, we will point out that in countries that enforce laws based on the strict product liability doctrine, those who market AI systems as goods must face greater liabilities than those who offer services in the subject.[32]

Not minimizing the extra-contractual liabilities we have mentioned, the experience as to computer software shows us that most of the problems posed in the practice are related to contractual liabilities. That links this item with the item above, and urges to stress the relevance and recurrent frequency of the stipulations to limit the liability before the user and third parties.[3]

8. Conclusions

The study of intellectual property problems linked with AI, serves to stress the existence of new kinds of immaterial assets that result from great investments of human creativity, capital and business organization, which are extremely useful and valuable, but that (a) receive from the current law institutes a protection that does not cover all their components; and (b) result in easy appropriation by third parties. Everything seems to indicate that the solution to these problems can be efficiently given within the traditional categories of intellectual property. But, undoubtedly, we must apply them in a different way and provide many times specific solutions.

Instead of causing degradation or ruin, the introduction of new products of intellectual activity and new problems has produced a beneficial crisis of growth and refinement in intellectual property law. We can anticipate that AI will be one of the factors that will motorize this process in the near future.

NOTES

[1] Following AI's processes, systems of image recognition compare "pixel by pixel" a scanned signature with a field model. Other systems, following sequencial methods, compare "byte per byte" an alphanumeric expression with others filed in a dictionary. Why should we attribute "intelligence" to the first systems and deny the same attribution to the second one?

[2] Robert Maniquant "Points de vue sur l'intelligence artificielle," EXPERTISES, No. 96, p. 212.

[3] From $66 million in 1983 to $8,5 billion in 1993, according to MICRO ELECTRONIC MONITOR, Apr.-Sep. 84.

[4] Source: Club Intelligence Artificielle (France). Based on the data projection as of March 30, 1986.

[5] See "Systèmes experts et Droit," directed by Marc Schauss, CRID, Namur 1988; and "Los Sistemas Expertos y la Ley," by Jaap H. Spoor, in DAT, No. 26, October 1990, page 1.

[6] This kind of system covers most of AI's resources included in the equipment of navigation, combat, etc.

[7] Note that this category does not include packaged systems that require from the user only the indication of certain parameters indispensable to direct the "intelligent" process to its objective.

[8] In fact, I have not learnt that such a product is being currently offered for AI systems, but as this is a common form of marketing works in traditional software systems and data bases, it seems logical that they appear at any moment.

[9] In relation to the legal "multiprotection" that favors computer software, I pointed out in a previous study, that the possibility of superposing Patent Law and Copyright protection is also present in works of other nature, such as architectural and scenographic ("Los más recientes casos norteamercianos en materia de Propiedad Intelectual sobre Software," DAT, No 28/29).

[10] I resort to the quotation marks given my uncertainty about the legal nature of immaterial creations of "form" or "essence" that do not have a human inventor or author.

[11] The patent for "discoveries" protect the exclusive use of the industrial application of an element or force of nature, but not the exclusive use of such category of elements or forces. The possession of a geologic or botanic piece, of a special decorative value or beauty, makes us the owners of movable property, but it does not agree on a title for the exclusive exploitation of its "form" as to an immaterial asset.

[12] Michel Vivant and Christian Le Stanc, "Lamy, droit de l'informatique 90," Paris 1990, page 816, warns against the risks of the contemporary "phenomenon of the Intellectual Property atomizations."

[13] I greatly appreciate the capacity of intellectual property rights--in particular, copyright law--to interpret technical facts in constant change on the basis of permanent juridical regulations. Thanks to that, television, satellites, computers, etc., neither produced a fragmentation of the institution nor a "regulatory pollution," equally dangerous. Therefore, I express my preference towards "unitary visions" of intellectual property, such as Ascarelli's or Troller's.

[14] See, for instance, the cases quoted in "Patentable Subject Matter: Mathematical Algorithms and Computer Programs," by Lee E. Barret, in SOFTWARE PROTECTION, October 1989, page 6.

[15] This is the doctrine that prevails in the most recent American software clone cases, explained in detail in Lotus Development Corporation v. Paperback Software International, 6.28.90, Civil No. 87-76-K, U.S.D.C. Massachusetts.

[16] Claude Colombet prefers and justifies the uncertainty generated by a "rather inexact notion the Courts must appreciate in each case," arguing that a "strict definition (of the protected works), in the case it could be possible, would have presented more inconveniences than advantages: not being able to represent a perfect hypothesis, it would have confined the judges to a limited and closed field" ("Propriété Littéraire et Artistique," Paris 1990, page 39).

[17] I make use of the, in my opinion, very graphic and fortunate expression chosen to explain the difficult action of separating the idea from the expression, by Justice Keeton in his remarkable opinion in "Lotus Development Corp. v. Paperback Software International" (see quote 15).

[18] See "Le droit des savants," by Paul Olagnier, Paris 1937.

[19] See "Neural Networks: The Next Intellectual Property Nightmare?," by Andy Johnson Laird and "Protection of Intellectual Property in Neural Networks," by Gerald H. Robinson, both in THE COMPUTER LAWYER, March 1990, pages 7 and 17, respectively.

[20] In the matter of neural networks, if we tried to solve the problem of protection considering that the expression capable of receiving protection is that whose "form" appears at the moment of the "input," we will topple with the problem that such inputs are elements of such simplicity that they mingle with the idea itself.

[21] Briefness is not a negative condition of copyright protection, not only because the law protects the creation "of any nature and extension" (wording of Argentine Copyright Act, Section 1) but also because it is precisely the briefness that makes more valuable and original the expression of works of different sorts, such as the documental or the journalistic.

[22] To my knowledge, there is no legal experience, except in reference to certain conflicts of plagiarism, related to software applied to the "inference engine" function.

[23] According to André Bertrand, the production of the "inference engine" itself--recognized as a mere computer program by most of the people who considered the problem--will also be a collaboration hypothesis in most cases; given that both experts and programmers intervene in its creation because of its typical characteristics ("L'intelligence artificielle, la robotique, les systèmes experts et le droit," EXPERTISES, No. 96, p. 219).

[24] Robert Maniquant "Points de vue sur l'intelligence artificielle," EXPERTISES, No. 96, p. 217.

[25] That is how it was suggested by Michel Vivant and Christian Le Stanc, op. cit., p. 816.

[26] The question is not extreme in countries like Argentina, where the lack of registration of published works does not prevent the exercise of copyright protection actions. As to USA, the painful experience that results from the recent decision in "Ashton-Tate Corporation vs. Fox Software, Inc. et al" (12.11.90, Civil No. 88-6837 TJH (Tx) CD Cal.) makes it unnecessary to formulate any further considerations.

[27] Conf. Vivant and Le Stanc, op. cit., p. 817.

[28] See "Derecho de Autor y responsabilidad"

[29] The human emulation by AI systems has motivated not only worries about the liability, but also speculation of the possiblity of granting some kind of legal guardianship to the 'intelligent machines,' given their "subjectivity." See "Acerca de la tutela legal de los autómatas implementados por programas inteligentes," by Giancarlo Taddei Elmi, in DAT, No. 9, May 1989, page 1.

[30] Michael Kirby in "Legal and Ethical Issues in Artificial Intelligence," reminds us of the possibility that this occurs when the systems constitute the interface between a machine of medical use and the human life (INTERNATIONAL COMPUTER LAW ADVISER, Nov. 87, page 6).

[31] Conf. Zoppini, Andrea "Commercializzacione dei sistemi esperti e responsabilita civile". 40, Congresso informatica en Regolamentazioni Giuridiche, Roma 1988, III, 22, page 10.

[32] For American law, see "Expert Systems: Strategies for minimizing Liability" by Howard G. Zaharoff, THE COMPUTER LAWYER, February 1989, page 30. for European Community law, see "Errores en Software: una cuestion de vida y responsabilidad," by Guy P.V. Vandenberghe, in DAT, No. 3, November 1988, page 5.

[33] Zaharoff suggests some efficient contractual resources to reduce the liability bbefore the user, such as: avoid defining the system as "expert," require concrete acknowledgement of certain conditions and preventions from the user, require conservative actions of the user's right, etc. (op. cit, page 33).

THE CREATION OF WORKS OF COPYRIGHT UNDER JAPANESE COPYRIGHT LAW RESULTING FROM THE UTILIZATION OF ARTIFICIAL INTELLIGENCE

by

Shigeru Miki
Miki Law & Patent Office
Tokyo, Japan

I. Issues

As the development and diffusion of both computer hardware and software continues to occur at a rapid pace, along with the use of computers in all sectors of contemporary society, the use of computers to create various copyrightable works is, accordingly, growing.

There are various reasons for using a computer in the creative process, such as saving time and labor, or creating specialized forms of expression. Computer programs are being used in a wide variety of fields to create works of authorship such as music, art works, schematic drawings and diagrams, movies, programs, and even translations.

Specifically, computer generated works of authorship include such items as expert systems, CAD/CAM, national language systems, computer generated compositions, automatic programing tools for development of software, and support programs for creation of databases.

The main legal issue related hereto is whether works of authorship created by way of computer programs, or computer generated works, express sufficient human ideas and sentiments to deem them copyrightable.

Furthermore, numerous problems exist in relation to the granting of copyrights in these instances, such as the assessment from a legal perspective of which persons should be considered direct participants in the composition process, or who is the composer of the program or contributor to the composition of the database.

II. Copyright Author

First of all, I wish to explain the meaning of the terms "copyright author" and "intellectual creativity" under the current Japanese Copyright Law in order to resolve the issues present here. "Author," according to traditional copyright theory, denotes the individual who creates a copyrightable work. Even if such person is assisted by another, he still will be considered the author so long as he possesses the authorization to direct the assistant to prepare, change, or revise a work due to an employment agreement between them, or other development agreement. In the case of a corporate author, there are two legal requirements. First, a corporation must have established a specific project to create a particular work and second, such corporation should have the legal authority to direct the employees or individuals to create said work. Such employees or individuals must also create the work in the course of their employment or duties pursuant to the company's orders. Although a corporation does not actually create the work itself, it may direct employees or individuals to physically carry out such processes. Based upon this interpretation, it seems to me that an author need not physically create a work, but should participate in creation of the work. For example, in case of a computer-generated work, whether or not an end user becomes the author of such work depends on how he participated in the creation of such work. If his participation could be evaluated as an intellectual creative activity under the Japanese Copyright Law, he would qualify as an author of such work.

III. Intellectual Creativity under the Japanese Copyright Law

The Japanese Copyright Law protects works of authorship in which ideas or sentiments are expressed in a creative manner within the literary, scientific, artistic or musical domains (Article 2, Paragraph 1(1). An interim report by the No. 6 Sub-Committee of the Agency for Culture Affairs has stated that the intellectual creativity of computer programs stems from the creative combination of instructions and statements expressed therein (Interim Report, January, 1984). Japanese courts take the same position.

Then, the core question is whether a user participates in making the combination of instructions and statements expressed therein, or how the user would carry out development work of such a computer program.

IV. National Language System

Based upon those requirements, national language systems can be created which are copyrightable. Translated works generated by a Japanese language translation system are not, to date, sufficiently accurate. As a result, a user must review such translations and revise the wording, grammatical structure or style thereof. Therefore, such user may be considered the author of the final version of the translation.

V. Code Generator

A Code Generator includes a common tool used in assembling programs. An assembler program can be used to convert from the source code of an application program to an object code thereof. The object code generated by such assembler program is considered a reproduction of the source code so long as any single word embodied in the source code reflects the wording in the object code.

Tokyo district courts and other Japanese district courts have held in several cases that there is no intellectual creativity in converting one language to another. Therefore, such object code is only a copy of the source code. In a case of disassembly from the object code to source code, the courts have taken the same position.

VI. Expert Systems

An expert system consists of a knowledge base and inference engine. Whenever someone uses such expert system he obtains a particular solution or answer on the screen. The question that arises under law concerns who shall be deemed the author of the solution which appeared on the screen by way of using the expert system in a certain manner.

VII. Computer Generated Works Using Expert Systems

A variety of expert systems exist, some of which are intended to enable an end user to change, modify, or complete a work by himself. For example, a certain expert system can support the development of an application program. Based upon the data and optionized item, the expert system generates the application program desired. In this computer generated program, the computer contains many program components or "parts". The issue here relates to who is to be deemed the author of such application program. It can be stated that the developer of such expert system may be entitled to rights of authorship because any user of such expert system must choose a particular way to develop an application program within the scope of optimization selected. Thus, the developer controls the process of creation and can imagine what components of the computer generated program will ensue, including the structure thereof. Thus, end users may not create any application program within the scope of optimization.

There is a second type of expert system through which the user may put data into a computer and develop his own application program. In that case, the developer would not participate in developing such application program. Rather, the user must actively participate in the development of the program by selecting data, arranging the structure and revising the original framework. Based upon his participation, the user would be considered the author of the application program. However, such application program would be a derivative work of an original application program contemplated by the overall developer of the expert system.

There is a third type of expert system through which the user may put data into a computer. Such expert system would create a certain computer program without any assistance of human beings. Therefore, neither the developer of such expert system nor the user will participate in the creation of a particular work, by using such expert system. I have not yet heard of the actual development of such a computer system. However, when such an expert system is developed, we should make a decision about how a computer program generated by such expert system should be protected under intellectual property laws.

GENERAL REPORT

by

Paul Goldstein
Stella W. and Ira S. Lillick Professor Law
Stanford University
Palo Alto, California
United States of America

The Reporter congratulated WIPO for initiating this Worldwide Symposium on the Intellectual Property Aspects of Artificial Intelligence. He observed that the Symposium is innovative not only in being the first such WIPO Symposium to be held outside Geneva, but also in gathering together participants from such varied backgrounds as the public and private sectors, academic and practice, and law, science, and engineering. The discussion during the two and one-half days of the Symposium opened a window onto the exhilarating future that artificial intelligence holds out as a promise. The discussion in the Symposium also opened a window onto the brief history of intellectual property protection for computer programs.

Several participants in the Symposium lamented the prevailing confusion in the general area of intellectual property protection for computer-related products. Comments from other participants suggested, however, that the source of this lament is not so much confusion as complexity and multiplicity--complexity in the underlying technology and multiplicity in the legal viewpoints on proper means of protection. Some legal viewpoints were conservative--captured perhaps in the slogan, "If it's not broke, don't fix it." Other view points expressed at the Symposium were more ambitious, reflected in the title of Professor Davis' talk, "The Assumptions Are Broken." A pragmatic legal view also emerged at the Symposium, centering on Director General Bogsch's query respecting protection for artificial intelligence, and products resulting from artificial intelligence technologies: "What are the acts that require authorization from whom?"

Four questions underlie most if not all of the presentations and interventions at the Symposium:

1. What aspects of artificial intelligence should be the object of study?

2. What are the requisite conditions for creativity in artificial intelligence?

3. What institutions, legal or otherwise, will enhance the conditions for creativity in artificial intelligence?

4. What is the appropriate method for inquiring into the legal institutions affecting creativity in artificial intelligence?

1. What aspects of artificial intelligence should be the object of study?

Some participants in the Symposium viewed artificial intelligence through a narrow legal lens, and characterized it as no more than a combination of two elements that already occupy a defined place in intellectual property law: computer programs and databases. Viewed more capaciously, artificial intelligence is, as Professor Winograd observed, a teleology, moving toward machines that have the power to think. (In this sense, the study of artificial intelligence also implicates the study of human intelligence.) Viewed pragmatically, artificial intelligence is, in Mr. Bhojwani's conception, the ability to produce knowledge, a tool for creating products and products included in other technologies.

Most commentators at the Symposium who addressed the question agreed that artificial intelligence centers on behavior, and that the particular form of material support in which behavior is embodied--hardware or software--has no relevance for purposes of science or policy (although it may presently have significance for applicable legal regimes). Software, Professor Davis noted, "is a machine whose medium is text."

Artificial intelligence is digital--a characteristic that enhances its replicability, distribution and multiple usage. Professor Miller questioned whether this characteristic implies a difference from the past, for such widely distributed phenomena as concert performances have occurred for ages. But Professor Davis' presentation suggested that digitalization in fact implies a shift in paradigms. For the first time, users will have access to vast networks through which they can manipulate their information environment and themselves become artists and communicators--on a scale as grand as the globe or as intimate as Professor Rucker's cyberspace.

Several comments suggested that, a short time from now, the wondrous examples of artificial intelligence depicted at this Symposium might look more like antiquated fossils. Other comments suggested, if the potential of artificial intelligence is to be realized, it is important today for public policy to promote, not discourage, creativity.

2. What are the requisite conditions for creativity in artificial intelligence?

Professor Winograd and Professor Nagao agreed that it will require major scientific revolutions for artificial intelligence to realize its promise. An offhand observation by Professor Rucker respecting "cyberspace" suggests the nature of law's role in fostering--or impeding--scientific development: "This is so new an area that we haven't encountered any intellectual property problems yet." Should intellectual property be a problem, or should it be a solution? Professor McCarty underscored that intellectual property should not stand in the way of scientific progress.

Professor Reichman observed that the paradox of the conservative legal approach to intellectual property in this area is that, by preserving existing legal categories, it may in fact be destabilizing the conditions of creativity. Intellectual property implies not only incentives to create, but freedom to create. The danger of existing legal categories is that, constructed with yesterday's technologies in view, they may impair the building and borrowing that is essential to creative science and engineering.

Databases offer an example. Mr. Weyer asked: Will legal rules introduce a renaissance or only an administrative and legal overload? Will otherwise fluent communication channels fall victim to traffic jams and information gridlock? (By coincidence, on the last day of the Symposium, the Supreme Court of the United States delivered a decision having important implications for intellectual property in databases, holding in Feist Publications, Inc. v. Rural Telephone Service Co., that alphabetized listings in telephone directory white pages are not copyrightable and that so-called "sweat of the brow" does not qualify for copyright protection.) In the larger context of electronic database networks, Mr. Akin Thomas observed the importance of attending to the divide between economically developed and economically developing countries; the great promise of these technologies is to bridge the widening gap between the world's haves and have-nots.

3. What institutions, legal or otherwise, will enhance the conditions for creativity in artificial intelligence?

The great bulk of artificial intelligence research to date has been supported not by intellectual property's promise of marketplace rewards, but rather by government subsidy. Some commentators at the Symposium suggested that, in this quarter, intellectual property will do best if it simply stays out of the way of scientific research. But, with respect to applied science, no one at the Symposium expressed any doubt that the market mechanism of intellectual property would represent the dominant inducement to investment in the coming years.

4. What is the appropriate method for inquiring into the legal institutions affecting creativity in artificial intelligence?

How are policymakers to go about the task of designing intellectual property systems to encourage the appropriate level of private investment in artificial intelligence? As evidenced by many comments at the Symposium, the way in which policymakers frame their inquiry will control the conclusions they reach. Three approaches to legal inquiry emerged at the Symposium.

First is the approach characterized by the comments of Mr. Goldberg, Professor Miller and, to some extent, Professor Spoor. Essentially, this approach reasons by analogy and by category: Artificial intelligence systems consist of computer programs and of databases; computer programs are like literary works and thus are copyrightable; databases are compilations and thus are copyrightable; q.e.d., artificial intelligence is copyrightable. Even though computer-generated products do not in fact bear the impress of an author's hand they, at least superficially, look like the traditional works contemplated by both the copyright and author's right traditions. Since these products look like the works of Picasso and sound like the works of Verdi they should--this approach maintains--be protectable under both copyright and author's right regimes.

Some participants at the Symposium challenged this approach as indeterminate and fundamentally incoherent from a research and development policy perspective, and stated that it could ultimately disserve creativity in artificial intelligence and allied fields. To rely on existing intellectual property systems will attract investment to some forms of artificial intelligence, but not other, possibly more desirable, technologies that do not fit the copyright or patent models. To use Mr. Johnson-Laird's metaphor, certain islands in the software swamp may go unvisited.

A second approach to legal inquiry expressed at the Symposium, would begin with Mr. Bhojwani's question: What human values does the public policy seek to serve? At an intermediate level, this approach views technological behavior in the round, asking what intellectual property regimes are needed to attract optimal investment in artifical intelligence. As Professor Davis observed, sui generis protection does not necessarily imply specialized protection. It may, in fact, mean capturing the heart and soul of what is pervasively creative in today's and tomorrow's inventive endeavor.

A third method of legal inquiry expressed at the Symposium is the pragmatic approach captured in Dr. Bogsch's question: "What are the acts that require the authorization of whom?" One possible implication of this question is that if an existing intellectual property regime--copyright, for example--cannot answer these questions plausibly and coherently in the case of artificial intelligence, and the products created through artificial intelligence, that intellectual property regime may be inappropriate to the object at hand.

Professor Davis displayed several computer-generated products in the form of drawings that bore a surface resemblance to traditional works of art having human authors. Will the copying of these products require consent? Whose consent? The consent of the author of the original, generative program? Of the computer? Of the end user? Paradoxically, these products might more easily find protection under the author's right tradition than under the Anglo-American copyright tradition. Some impress, however faint, of the author's hand can be traced from the end product to the generative artificial intelligence program. But the utilitarian tradition of Anglo-American copyright law, which focuses on the economic incentives needed to produce the desired economic result, may find that payment of copyright tribute in perpetuity for the millions of drawings produced by a program at the cost only of electricity is not needed as an incentive to produce the programs that generate these works.

Discussion at the Symposium also indicated the importance of attending to external legal constraints. Mr. Goldberg observed, for example, that a copyright approach to computer programs has become embedded in the laws of several nations, in the WIPO Model Law and in the proposed EC Directive. However, Professor Reichman suggested, in his example of the turn-of-the-century experience with design protection, that these are not natural laws like the law of gravity, but are human laws that humans can alter.

Several interventions at the last session of the Symposium suggested that, as artificial intelligence technologies continue to evolve, it will be useful regularly to monitor their intellectual property aspects to determine whether existing intellectual property regimes adequately foster creativity in the area. One participant observed that such oversight activities had a logical place on WIPO's overall agenda, and another participant proposed that meetings like the present Symposium be sponsored by WIPO in the future.

LIST OF PARTICIPANTS

I. SPEAKERS

Norman ALTERMAN, Senior Vice President, Legal Affairs, Motion Picture Export Association of America (MPEAA), New York, New York, United States of America

Robert BARR, Partner, Townsend & Townsend, Palo Alto, California, United States of America

Ashok BHOJWANI, Managing Director, TSG Consultants Limited, New Delhi, India

Randall DAVIS, Associate Director, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America

Thomas K. DREIER, Head, French Law Department, Max Planck Institute, Munich, Germany

Morton D. GOLDBERG, Attorney-at-Law, Schwab Goldberg Price and Dannay, New York, New York, United States of America

Paul GOLDSTEIN, Professor of Law, Stanford University, Stanford, California, United States of America

Péter GYERTYÁNFY, Legal Director, Hungarian Bureau for the Protection of Author's Rights (ARTISJUS), Lecturer, Eötvös Lorand University, Budapest, Hungary

Andy JOHNSON-LAIRD, President, Johnson-Laird Inc., Portland, Oregon, United States of America

Ronald S. LAURIE, Attorney-at-Law, Irell and Manella, Menlo Park, California, United States of America

L. Thorne McCARTY, Professor, Department of Computer Sciences, Rutgers University, New Brunswick, New Jersey, United States of America

Shigeru MIKI, Lecturer, Chuo University, Senior Partner, Miki Law and Patent Office, Tokyo, Japan

Antonio MILLE, President, Latin America Institute for Advanced Technology, Computer Science and Law, Buenos Aires, Argentina

Arthur R. MILLER, Professor, Harvard Law School, Cambridge, Massachusetts, United States of America

Makoto NAGAO, Professor, Kyoto University, Kyoto, Japan

Jerome REICHMAN, Professor, Vanderbuilt University, Nashville, Tennessee, United States of America

Rudolph v. B. RUCKER, Professor, San Jose State University and Autodesk, Inc., Sausalito, California, United States of America

Daniel M. RUSSELL, Systems Sciences Laboratory, Xerox Palo Alto Research Center (PARC), Palo Alto, California, United States of America

Jaap H. SPOOR, Partner, Trenité van Doorne, Professor, Vrije Universiteit, Amsterdam, Netherlands

Clifford M. STURT, Partner, Marks & Clark, London, United Kingdom

Akin THOMAS, Managing Director, Heinemann Educational Books, Lagos, Nigeria

Stephen A. WEYER, Manager, Intelligent Applications, Apple Computer, Inc., Cupertino, California, United States of America

Terry WINOGRAD, Professor, Stanford University, Stanford, California, United States of America

II. STATES

AUSTRALIA

Peter TREYDE, Secretary, Copyright Law Committee, Business Division, Attorney-General's Department, Barton, ACT

CAMEROON

Hyacinthe FOUDA, Director of Legal Studies, Minister of Information and Culture, Yaoundé

CHINA

Gao HANG, Deputy Chief, International Division, National Copyright Administration of China (NCAC), Beijing

COSTA RICA

Andrés HERNANDES OSTI, Registro de la Propiedad Industrial, Ministerio de Justicia, San Jose

FINLAND

Rainer OESCH, Department of Private Law, University of Helsinki, Helsinki

HUNGARY

Péter GYERTYÁNFY, Legal Director, Hungarian Bureau for the Protection of Author's Rights (ARTISJUS), Budapest

ITALY

Marco FORTINI, Ambassador, Head, Intellectual Property Office, Ministry of Foreign Affairs, Rome

JAPAN

Takayuki SUMITA, Assistant Director, Information Promotion Division, Machinery and Information Industries Bureau, Ministry of International Trade and Industry, Tokyo

KENYA

L.O. SESE, Deputy Director, Technical Department, Kenya Industrial Property Office, Nairobi

THAILAND

Somkid CHOTIGAVANIT, Deputy Director General, Fine Arts Department, Ministry of Education, Bangkok

UNITED STATES OF AMERICA

Phill GILL, Head, Literary Section II, Examining Division, Copyright Office, Washington, D.C.

Julia HUFF, Head, Performing Arts Section, Examining Division, Copyright Office, Washington, D.C.

Clifton KNIGHT, Copyright Office, Washington, D.C.

Nancy LAWRENCE, Head, Literary Section I, Examining Division, Copyright Office, Washington, D.C.

Susan O. MANN, Attorney-Advisor, U.S. Patent and Trademark Office, Washington, D.C.

Harriet OLER, Chief, Examining Section, Copyright Office, Washington, D.C.

Ralph OMAN, Register of Copyrights, Copyright Office, Washington, D.C.

William PATRY, Policy Planning Advisor to the Register of Copyrights, Copyright Office, Washington, D.C.

Marybeth PETERS, Policy Planning Adviser to the Register of Copyrights, Copyright Office, Washington, D.C.

Dorothy SCHRADER, General Counsel and Associate Register for Legal Affairs, Copyright Office, Washington, D.C.

Harvey WINTER, U.S. Department of State, Washington, D.C.

III. INTERGOVERNMENTAL ORGANIZATIONS

UNITED NATIONS DEVELOPMENT PROGRAMME (UNDP)

Carlos CORREA, Consultant, Regional Coordinator, Regional Programme on Co-operation in Informatics and Microelectronics in Latin America and the Caribbean, Buenos Aires, Argentina

COMMISSION OF THE EUROPEAN COMMUNITIES

Bridget CZARNOTA, Cosultant, Directorate General III, Internal Market and Industrial Affairs, Brussels, Belgium

IV. NON-GOVERNMENTAL ORGANIZATIONS

Committee of National Institutes of Patent Agents (CNIPA): Jürgen BETTEN (Patent Attorney)

International Association for the Protection of Industrial Property (AIPPI): Ronald S. LAURIE (Member, Executive Committee of the United States Group).

International Bureau of Societies Administering the Rights of Mechanical Recording and Reproduction (BIEMN): Ralph OMAN (Member, Legal and Legislation Committee of CISAC)

International Chamber of Commerce (ICC): Morton David GOLDBERG (Copyright Subcommittee Chair, U.S. Council for International Business)

International Confederation of Societies of Authors and Composers (CISAC): Ralph OMAN (Member, Legal and Legislation Committee)

International Copyright Society (INTERGU): Stanley J. ROTHENBERG (Member of the Committee)

International Federation of Industrial Property Attorneys (FICPI): John ORCHARD (President-UK), and Clifford Mark STURT (Marks and Clerk, London)

International Federation of Film Producers Associations (FIAPF): Norman ALTERMAN (Counsellor)

International Publishers Association (IPA): Joseph Alexis KOUTCHOUMOW (Secretary General)

Licensing Executives Society (International) (LES): Luis MEJIA (Representative; Licensing Associate, Stanford University)

Union of European Practitioners in Industrial Property (UEPIP): Jürgen BETTEN (President, Software Commission)

V. OTHER PARTICIPANTS

John BARTON, Stanford University, Stanford, California, United States of America

David CARSON, Partner, Schwab Goldberg Price and Danny, New York, New York, United States of America

Peter CHEUNG, Law Reform Commission, Hong Kong, China

Peter CHOY, Sun Microsystems, Inc., Sunnyvale, California, United States of America

Joyce CUTLER, Bureau of National Affairs, Washington, D.C., United States of America

Robert DEUTSCH, Hoover Institution, Stanford, California, United States of America

Seth GREENSTEIN, McDermott, Will & Emery, Washington, D.C., United States of America

Robert J. HART, Intellectual Property International Limited, Liverpool, United Kingdom

Marcelo HOFFMAN, SRI International, Menlo Park, California, United States of America

Steven HOWARTH, Hitachi America, Ltd., Brisbane, California, United States of America

Taisuke KATO, Toshiba America, Inc., Irvine, California, United States of America

Masanobu KATOH, Fujitsu Limited, Washington, D.C., United States of America

Jerry KLEIN, The Chatham Group, Los Altos, California, United States of America

Naoki KYOMOTO, NEC Corporation, Tokyo, Japan

Patrice LYONS, Attorney-at-Law, Washington, D.C., United States of America

John MOSS, Stanford University, Stanford, California, United States of America

Bernard NONNENMACHER, Stanford University, Stanford, California, United States of America

Susan NYCUM, Baker & McKenzie, San Francisco, California, United States of America

Kiichi OGURA, Mitsubishi Electric, Washington, D.C., United States of America

J.R. OKONESKI, Foster City, California, United States of America

Mirian PITTA, IBM Corporation, Botafogo, Rio de Janeiro, Brazil

Tom SCHWARTZ, Electronic Engineering Times, Mountain View, California, United States of America

Victor SIBER, Intellectual Property Law Department, IBM International Business Machines Corporation, Purchase, New York, United States of America

Jim SIMPSON, SLAC, Stanford, California, United States of America

Katherine SPELMAN, Townsend and Townsend, San Francisco, California, United States of America

Robert SPILLERS, IBM Corporation, San Jose, California, United States of America

Calvin WARD, McCubbrey, Bartels, Meyer & Ward, San Francisco, California, United States of America

VI. WORLD INTELLECTUAL PROPERTY ORGANIZATION (WIPO)

Arpad BOGSCH, Director General

Paul CLAUS, Director, Classifications and Industrial Property

Mihály FICSOR, Director, Copyright Law Division

Francis GURRY, Special Assistant, Office of the Director General

Jean ANTICEVIC, Liaison Officer

- * - * -