

The Origins of the American Association for Artificial Intelligence (AAAI)

Raj Reddy

■ This article provides a historical background on how AAAI came into existence. It provides a rationale for why we needed our own society. It provides a list of the founding members of the community that came together to establish AAAI. Starting a new society comes with a whole range of issues and problems: What will it be called? How will it be financed? Who will run the society? What kind of activities will it engage in? and so on. This article provides a brief description of the considerations that went into making the final choices. It also provides a description of the historic first AAAI conference and the people that made it happen.

The Background and the Context

While the 1950s and 1960s were an active period for research in AI, there were no organized mechanisms for the members of the community to get together and share ideas and accomplishments. By the early 1960s there were several active research groups in AI, including those at Carnegie Mellon University (CMU), the Massachusetts Institute of Technology (MIT), Stanford University, Stanford Research Institute (later SRI International), and a little later the University of Southern California Information Sciences Institute (USC-ISI).

My own involvement in AI began in 1963, when I joined Stanford as a graduate student working with John McCarthy. After completing my Ph.D. in 1966, I joined the faculty at Stanford as an assistant professor and stayed there until 1969 when I left to join Allen Newell and Herb Simon at Carnegie Mellon University (CMU). The 1960s at Stanford AI Labs (SAIL)



Raj Reddy.

A Sampling of AI Research in 1980, from the Proceedings of AAAI-80

Vision

Early Vision Processing

Recovering Surface Orientation from Texture, *Andrew P. Witkin, SRI International*

Shape-from-Texture Paradigm, *John R. Kender and Takeo Kanade, Carnegie Mellon University*

Low Level Vision Systems, *William B. Thompson and Albert Yonas, University of Minnesota*

Interpreting Line Drawings as 3D Surfaces, *Harry G. Barrow and Jay M. Tenenbaum, SRI*

Shape Encoding and Subjective Contours, *Brady, Grimson, MIT and Langridge, CSIRO*

Scene Analysis

Information Needed to Label a Scene, *Eugene C. Freuder, University of New Hampshire*

Interpretive Vision and Restriction Graphs, *Rodney Brooks and Thomas Binford, Stanford*

Sticks, Plates, and Blobs: A 3D Object Representation for Scene Analysis, *Shapiro, Moriarty, Mulgaonkar, and Haralick, Virginia Polytechnic Institute and State University*

Motion Analysis

Constraint-Based Inference from Image Motion, *Daryl T. Lawton, University of Massachusetts*

Static Analysis of Moving Jointed Objects, *Jon A. Webb, University of Texas at Austin*

Bootstrap Stereo, *Marsha Jo Hannah, Lockheed Palo Alto Research Laboratory*

Robotic Vision

Locating Partially Visible Objects: The Local Feature Focus Method, *Robert C. Bolles, SRI International*

Collision Avoidance Among 3D Objects, *Ahuja, Chien, Yen, and Bridwell, University of Illinois*

Automated Inspection Using Gray-Scale Statistics, *Stephen T. Barnard, SRI International*

Human Movement Understanding,

Badler, O'Rourke, Platt, and Morris, University of Pennsylvania

Using Contextual Information in Computer Vision, *Olivier D. Faugeras, USC*

Program Synthesis

Question Ordering in Mixed Initiative Program Specification Dialogue, *Louis Steinberg, Rutgers*

Some Algorithm Design Methods, *Steve Tappel, Systems Control, Inc.*

Automatic Goal-Directed Program Transformation, *Stephen Fickas, USC ISI*

Incremental, Informal Program Acquisition, *Brian McCune, AI and D Systems*

A Basis for a Theory of Program Synthesis, *P. A. Subrahmanyam, USC ISI*

A Program Model for Computer Aided Program Synthesis, *Richard J. Wood, University of Maryland*

Theorem Proving

An Efficient Relevance Criterion for Mechanical Theorem Proving, *David A. Plaisted, University of Illinois Urbana-Champaign*

On Proving Laws of the Algebra, *Jacek Leszczylowski, Polish Academy of Sciences*

Establishing Completeness Results in Theorem Proving, *Peterson, University of Missouri at St. Louis*

Automatic Generation of Semantic Attachments in FOL, *Luigla Aiello, Stanford*

HCPRVR: An Interpreter for Logic Programs, *Daniel Chester, University of Texas at Austin*

First Experiments with Rue Automated Deduction, *Vincent J. Digricoli, The Courant Institute*

Mathematical and Theoretical Foundations

What's Wrong with Non-Monotonic Logic? *David J. Israel, Bolt Beranek and Newman, Inc.*

Pathology on Game Trees: A Summary of Results, *Dana S. Nau, University of Maryland*

Max-Min Chaining of Weighted Causal Assertions Is Loop Free, *Ng and Walker, Rutgers University*

Problem Solving

Cooperating Expert Systems

HEARSAY-III: A Framework for Expert Systems, *Balzer, Erman, London, and Williams, USC ISI*

Quantifying and Simulating the Behavior of KBIS, *Lesser, Reed, and Pavlin, University of Massachusetts*

Representation of Task Knowledge in User Interfaces, *Eugene Ball and Phil Hayes, Carnegie Mellon University*

Problem Solving and Control

Representation of Control Knowledge in Expert Systems, *Janice S. Aikins, Stanford University*

DELTA-MIN: A Search-Control Method for Information-Gathering Problems, *Carbonell, Carnegie Mellon University*

On Waiting, *Arthur M. Farley, University of Oregon*

A Planner for Reasoning about Knowledge and Action, *Douglas E. Appelt, SRI International*

Making Judgments, *Hans J. Berliner, Carnegie Mellon University*

Multiple-Agent Planning Systems, *Kurt Konolige and Nils J. Nilsson, SRI*

A Simple Game-Searching Algorithm with Proven Optimal Properties, *Judea Pearl, UCLA*

Problem Solving in Using Interactive Dialog, *Harry C. Reinstein, IBM Palo Alto Scientific Center*

Representing Knowledge in an Interactive Planner, *Ann E. Robinson and David E. Wilkins, SRI*

Inference with Recursive Rules, *Stuart C. Shapiro and Donald P. McKay, SUNY Buffalo*

Knowledge Representation

Advanced Knowledge Representation

A Frame-Based Production System Architecture, *David E. Smith and Jan E. Clayton, Stanford*

Knowledge Embedding in the Description System Omega, *Hewitt, Attardi, and Simi, MIT*

A Representation Language, *Russell Greiner and Douglas B. Lenat, Stanford*

Applied Knowledge Representation

Spatial and Qualitative Aspects of Reasoning about Motion, *Kenneth D. Forbus, MIT*

Computer Interpretation of Human Stick Figures, *Martin Herman, Carnegie Mellon University*

Research on Expert Problem Solving in Physics, *Novak and Araya, University of Texas at Austin*

Knowledge-Based Simulation, *Philip Klahr and William S. Faught, The Rand Corporation*

Interactive Frame Instantiation, *Carl Engelman, Ethan A. Scarl, and Charles H. Berg, MITRE*

Specialized Issues in Knowledge Representation

Descriptions for a Programming Environment, *Ira Goldstein and Daniel Bobrow, Xerox PARC*

Rule-Based Inference in Large Knowledge Bases, *William Mark, USC ISI*

A Process for Evaluating Tree-Consistency, *John L. Goodson, Rutgers University*

Reasoning about Change in Knowledgeable Office Systems, *Gerald R. Barber, MIT*

On Supporting the Use of Procedures in Office Work, *Fikes and Henderson, Jr., Xerox PARC*

Metaphors and Models, *Michael R. Genesereth, Stanford University*

About Authority Structures You Were Unable to Represent, *James R. Meehan, University of California, Irvine*

Real Time Causal Monitors for Complex Physical Sites, *Rieger and Stanfill, University of Maryland*

Knowledge Acquisition

Applying General Induction Methods to the Card Game Eleusis, *Tom Dietterich, Stanford*

Modeling Student Acquisition of Problem-Solving Skills, *Robert Smith, Rutgers University*

A Computer Model of Child Language Learning, *Mallory Selfridge, Yale University*

The Instructable Production System Project, *Michael D. Rychener, Carnegie Mellon University*

Using a Matcher to Make an Expert

Consultation System Behave Intelligently, *René Reboh, SRI*

An Approach to Acquiring and Applying Knowledge, *Norman Haas and Gary G. Hendrix, SRI*

Self-Correcting Generalization, *Stephen B. Whltehill, University of California, Irvine*

Specialized Systems

Intelligent Retrieval Planning, *Jonathan J. King, Stanford University*

A Theory of Metric Spatial Inference, *Drew McDermott, Yale University*

Design Sketch for a Million-Element NETL Machine, *Scott E. Fahlman, Carnegie Mellon*

Perceptual Reasoning in a Hostile Environment, *Thomas D. Garvey and Martin A. Fischler, SRI*

Overview of an Example Generation System, *Edwina L. Rissland and Elliot M. Soloway, University of Massachusetts*

Structure Comparison and Semantic Interpretation of Differences, *Wellington Yu Chiu, USC ISI*

Performing Inferences over Recursive Data Bases, *Naqvi and Henschen, Northwestern University*

Piaget and Artificial Intelligence, *Jarrett K. Rosenberg, University of California, Berkeley*

Applications

R1: An Expert in the Computer Systems Domain, *John McDermott, Carnegie Mellon University*

Rule-Based Models of Legal Expertise, *D. A. Waterman and Mark Peterson, Rand Corporation*

Exploiting a Domain Model in Expert Spectral Analysis, *David Barstow, Schlumberger Research*

A System for the Automatic Analysis of Business Correspondence, *Lance Miller, IBM Research*

A Knowledge Based Design System for Digital Electronics, *Milton R. Grinberg, University of Maryland*

Theory Directed Reading Diagnosis, *Christian Wagner and John Vinsonhaler, Michigan State*

A Word-Finding Algorithm with a Dynamic Lexical-Semantic Memory for Patients with Anomia Using a Speech Prosthesis, *Colby, Christinaz, Graham, and Parkison, UCLA*

Trouble-Shooting by Plausible Inference, *Leonard Friedman, JPL, Caltech.*

An Application of the Prospector System to Uranium Resource Evaluation, *John Gaschnig, SRI*

Some Requirements for a Computer-Based Legal Consultant, *L. Thorne McCarty, Rutgers*

Natural Language

When Expectation Fails: A Self-Correcting Inference System, *Richard H. Granger, Jr., University of California Irvine*

Generating Relevant Explanations: Natural Language Responses to Questions about Database Structure, *Kathleen R. McKeown, University of Pennsylvania*

The Semantic Interpretation of Nominal Compounds, *Timothy Wilking Finin, University of Illinois*

Towards an AI Model of Argumentation, *Birnbaum, Flowers, and McGuire, Yale University*

Knowledge Representation for Syntactic / Semantic Processing, *Bobrow and Weber, University of Pennsylvania*

Language and Memory: Generalization as a Part of Understanding, *Michael Lebowitz, Yale*

Failures in Natural Language Systems: Applications to Data Base Query Systems, *Mays, University of Pennsylvania*

Memory Models

Organizing Memory and Keeping It Organized, *Janet L. Kolodner, Yale University*


Meta-Planning, *Robert Wilensky, University of California, Berkeley*

Narrative Text Summarization, *Wendy G. Lehnert, Yale University*



Stanford Artificial Intelligence Laboratory, 1968.

Photograph Courtesy Les Ezrnest.



Don Walker.

under the leadership of John McCarthy was indeed the golden age of AI at SAIL. Many of the areas of research at SAIL at that time continue to be active areas of research to this day. These include robotics, speech, vision, natural language processing, knowledge-based systems and heuristic programming, logic-based AI, and systems for playing games such as chess and checkers.

In spite of much active research in AI at the many centers around the country, there were no organized mechanisms for fostering the field. The first organized effort was the holding of the International Joint Conferences on AI (IJCAI) in 1969. This effort was organized by Rebecca Prather and Don Walker of SRI. This biennial conference was the only mechanism that existed throughout the 1970s for AI researchers to get together.

The 1977 IJCAI conference was held at MIT, and I was the program chairman that year. Organizing a program committee, reviewing the papers, creating a coherent technical program, and producing and transporting the proceedings turned out to be a nontrivial task for a group of faculty members whose primary responsibilities were teaching and research!

The 1979 IJCAI conference was held in Tokyo, and I was the general chairman tasked to deal with many of the organizational issues. Site selection, financial management, interna-

tional coordination, communicating with the global community, and identifying and cajoling group leaders for different tasks needed continual attention.

Why We Needed Our Own Society

The experiences of IJCAI-77 and IJCAI-79 convinced me that we need an enduring administrative mechanism and institutional memory such as those that are traditionally available at organizations like the Institute of Electrical and Electronic Engineers (IEEE) and the Association for Computing Machinery (ACM). For our community to thrive, it was clear that we needed more than just a biennial conference. Taking a cue from other scientific societies, it seemed desirable for AI to have its own society. We could have been part of one of the existing societies, but the time and the effort needed to get them to accept AI was too daunting.

To quote Allen Newell from his "AAAI President's Message" published in the first issue of *AI Magazine* (and reprinted in this issue):

Why did we come into existence? Note, I did not ask why we exist. Wherever, in our culture, a science is explored, a scientific society arises to nurture that exploration. AAAI exists because the science of artificial intelligence is being actively and vigorously explored in the United States.

... The one part of the birthday tale I haven't told is about the people. Societies are *sui generis*. They pass from nothing to something by an act of social will, in which particular people take to themselves the prerogative of forming themselves to be a scientific society for whatever (here, for artificial intelligence). Some set of people have to feel the calling and to make the inner decision at a particular point in history. For the AAAI the time was the recent IJCAI, held in Tokyo in August 1979. The people were almost entirely US participants on the IJCAI program and conference committees (including some who served prior IJCAIs)—a fact of mild significance, as I'll note in a moment. Their names appear as the founding council, whose formation by an act of self selection is required to get from nothing to something, society-wise. I was not one of them, so I can point out our (AAAI's) collective debt to them for being the founders.

It is necessary to take one more step in the narrowing social microdynamics of how AAAI got started. For it finally comes down to Raj Reddy, chairman of the IJCAI Board of Trustees and general chairman of the IJCAI in Tokyo, who took the initial personal act of decision that really started precipitation, and who carried all of us before him, until the AAAI was safely crystal-



I remember a memorable lunch at an unmemorable Chinese restaurant, where we sketched out the organizational design on a napkin, and Raj somehow talked me into agreeing to serve permanently as conference chair. I agreed, but only on the condition that I could hire Lou Robinson to do most of the work. So many other innovations came from Raj talking various people into things.

I also have some recollection of the first trade show we hosted at Stanford's Tressider Union. There were about 16 companies in a room that maybe was 100 by 100 feet at most: robots from companies like Machine Intelligence, natural language querying systems from companies like Artificial Intelligence Corp, and perhaps some early expert system shells. Within a few years, at our peak, we had taken over Cobo Hall, the main conference venue in Detroit!

—Marty Tenenbaum
(AAAI Conference Chairman, 1980–1987)

lized. The story of AAAI's origins comes safely to rest at this point, and need be spun out no further.

Founding members of the community that came together to establish the AAAI are given in table 1.

What Is the Name?

After many months of debates and discussion, we finally settled on the name of American Association for Artificial Intelligence and the acronym AAAI. Again to quote Allen Newell:

Officers and Chairmen

- President:* Allen Newell, Carnegie-Mellon University
- President-Elect:* Edward A. Feigenbaum, Stanford University
- Secretary-Treasure:* Donald E. Walker, SRI International
- Membership Chair:* Bruce G. Buchanan, Stanford University
- Publications Chair:* Lee D. Erman, USC/Information Sciences Institute
- AAAI Mag. Editor:* Alan M. Thompson, Jet Propulsion Laboratory
- Conference Chair:* Jay M. Tenenbaum, SRI International
- Program Chair:* Robert M. Balzer, USC/Information Sciences Institute
- Conf. Treasurer:* Lester D. Earnest, Stanford University
- Tutorial Chair:* Frederick Hayes-Roth, Rand Corporation

Executive Council

- Woody Bledsoe, University of Texas
- Bruce G. Buchanan, Stanford University
- Lee D. Erman, USC/Information Sciences Institute
- H. Penny Nii, Stanford University
- Nils J. Nilsson, SRI International
- D. Raj Reddy, Carnegie-Mellon University
- Earl D. Sacerdoti, SRI International
- Roger C. Schank, Yale University
- David L. Waltz, University of Illinois
- Patrick H. Winston, Massachusetts Institute of Technology

Table 1: Founding Organizers of AAAI.

AAAI-80 Program Committee

- Chairman:* Robert Balzer, University of Southern California
Information Sciences Institute
- Saul Amarel, Rutgers University
- Jon Bentley, Carnegie-Mellon University
- Woody Bledsoe, University of Texas
- Ron Brachman, Bolt Beranek & Newman
- Bruce Buchanan, Stanford University
- Edward Feigenbaum, Stanford University
- Ira Goldstein, Xerox PARC
- Peter Hart, SRI International
- Fredrick Hayes-Roth, Rand Corporation
- Carl Hewitt, Massachusetts Institute of Technology
- Doug Lenat, Stanford University
- David Luckham, Stanford University
- David Nitzan, SRI International
- Charles Rieger, University of Maryland
- Roger Schank, Yale University
- Jay M. Tenenbaum, SRI International
- David Waltz, University of Illinois

Table 2. AAAI-80 Program Committee.

The society has taken the name Artificial Intelligence. As all good AI'ers know, this name, introduced by John McCarthy in the fifties, has been controversial for quite awhile. It is often remarked that some of the controversy that swirls around our field is due to our name—the suspicion of people outside AI, including some in computer science, that our enterprise is not a legitimate scientific enterprise. ...I believe that the controversies have their natural cause in the type of knowledge our science reveals. They must be dealt with on the basis of substance and truth. So cherish the name Artificial Intelligence. It is a good name. Like all names of scientific fields, it will grow to become exactly what its field comes to mean.

Sustainable AAAI

Having decided that it may be desirable to have a society of our own, we needed to establish guiding principles for our society. The major concern was whether such a society could survive and become sustainable in the long term. It was agreed that each activity of the society would be self-sustaining. For example, we set the registration fees for the first AAAI conference around \$60 because the budgeting showed that we would break even at that price by holding the conferences at universities. Similarly the membership fee was supposed to just cover the cost of *AI Magazine* and one or two support staff members.

Any surpluses from year to year would go to an endowment. Annual income from the endowment would be used for outreach activities and to cover losses if any. We were fortunate in that, during the early years of our existence, AAAI had substantial surpluses primarily coming from the highly successful tutorial programs leading to an endowment of more than \$5 million. In the more recent lean years we have not been as successful. But by following the basic founding principles we've set forth, it is believed that we can be self-sustaining and stay alive for the foreseeable future.

Governance

We augmented the traditional structure of a president and council to enhance continuity and provide for smooth transition. The office of the president had three members: past president, president, and president-elect playing roles analogous to the roles of chairman, CEO, and COO. Staggered membership of the council also provided for preservation of continuity. All other appointments and terms were the responsibility of the president.

Activities of AAAI

The primary activities of AAAI are similar to those of other societies: conferences and workshops, journals and publications, membership and election of fellows. Much of the continuity from year to year was provided by the establishment of an AAAI administrative office in Menlo Park, California, which prospered under the able management and successive leadership of three executive directors: Lou Robinson, Claudia Mazzetti, and Carol Hamilton.

AI Magazine was created to help the community to stay abreast of significant new research and literature across the entire field of artificial intelligence. It was modeled after *Spectrum*, the IEEE magazine. Under the able successive leadership of Alan Thompson, Bob Englemore, Ramesh Patil, Elaine Rich, Jude Shavlik, and David Leake, serving as the editors in chief, *AI Magazine* has surpassed many of our initial expectations and has provided the all-important links to the community. A key ingredient for success was Mike Hamilton, who serves as the managing editor providing the all important corporate memory.

Perhaps the most satisfying aspect of embarking on this enterprise was the enthusiasm and excitement and camaraderie of the AI researchers at the first annual conference of AAAI at Stanford. Even after 25 years, it still stands out as one of the memorable conferences that I have ever attended. Most importantly it also provides a snapshot of AI research in the United States in 1980.

The First AAAI Conference

The broad spectrum of the research topics covered at the AAAI-80 conference was indicative of the vitality and vibrancy of our field at that time. There were multiple sessions in the areas of vision, problem solving, knowledge representation. In vision, we had topics such as early vision processing, scene analysis, motion analysis, and robotics vision. In problem solving, we had topics such as cooperating expert systems, problem-solving end control, and interactive problem solving. In knowledge representation, we had topics of advanced knowledge representation, applied knowledge representation, and special issues in knowledge representation. In addition we also had sessions on program synthesis, theorem proving, mathematical and theoretical foundations, knowledge acquisition, specialized systems, natural language processing, memory models, and applications of AI. Individual topics covered within these sessions are provided in the accompanying box. Surprisingly, in spite of ac-



I remember having a kind of distant amazement that the USA could find so much money in such a short time. AI societies in Europe were run on shoestrings, almost entirely by volunteer academics, but AAAI had piles of cash from day one (well, actually maybe something like day three) thanks to the AI explosion and the trade fair at the early meetings. Nothing like that had happened outside the USA.

—Pat Hayes (AAAI President, 1991–1993)

tive research programs in various centers, there were not many papers in the areas of learning, speech, and robotics.

The institutions represented at this conference were also interesting. Besides the usual AI powerhouses of CMU, MIT, Stanford, and SRI, we had multiple contributions from the University of Southern California Institute for Scientific Information, the University of Illinois at Urbana-Champaign, the University of Massachusetts, Yale University, the University of Pennsylvania, and Rutgers, reinforcing the premise that AAAI can be a forum to bring *all* members of the community together, avoiding the balkanization that can occur without a society of our own.

The authors of the papers at the first conference were also remarkable. They represented a Who's Who of AI, many of whom continue to be active to this day. What was unique about the attendees of the first conference was their



Allen Newell.

Photograph Courtesy Carnegie Mellon University.

genuine interest in the advancement of AI. Subsequently, during the AI bubble in the mid-1980s, a large number of people adopted AI as a quick way to become rich and famous. During the AI winter, many of these visitors disappeared, and now we seem to be back to the core group of researchers dedicated to the advancement of our field.

There were other activities at AAAI-80 that were equally exciting. There were a number of invited talks and panels from the leaders of our field on topics of immense interest to many of us. Unfortunately, we do not have a written record of their comments. Perhaps in the future, with advanced transcription and summarization technologies in hand, we will be able to keep a complete record of our views, agreements and disagreements, and expectations for the future. The program committee for the first conference, given in table 2, started with a clean slate and had to make up the rules as they went along. They deserve our respect and gratitude for the countless hours spent crafting an outstanding program.

Conclusion

The fact that we are celebrating the twenty-fifth year of the founding of AAAI is indeed

gratifying. It shows that our community has matured and has a widespread following and membership who care about the advancement of AI. Starting from our modest beginnings we have come a long way. A visit to our website shows a wide range of activities and functions we now provide for our membership

As we look to the future, we have a large unfinished agenda. The goal of creating artificial intelligences that reach or exceed human intelligence is still a long way away. We do not seem to have a coherent and coordinated vision and plan to create a human-level AI. I hope the next generation of AI researchers, using AAAI as a launching pad, can formulate a plan of action to get to our cherished goal.

Raj Reddy is the Mozah Bint Nasser University Professor of Computer Science and Robotics in the School of Computer Science at Carnegie Mellon University. He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences. He was president of the American Association for Artificial Intelligence from 1987 to 1989. Reddy was awarded the Legion of Honor by President Mitterand of France in 1984. He was awarded the ACM Turing Award in 1994. He served as cochair of the President's Information Technology Advisory Committee (PITAC) from 1999 to 2001.