

Conferences, Books, and Funding

ACCOMPANYING THE TECHNICAL PROGRESS IN ARTIFICIAL INTELLIGENCE DURING this period, new conferences and workshops were begun, textbooks were written, and financial support for basic research grew and then waned a bit.

The first large conference devoted exclusively to artificial intelligence was held in Washington, DC, in May 1969. Organized by Donald E. Walker (1928–1993) of the MITRE Corporation and Alistair Holden (1930–1999) of the University of Washington, it was called the International Joint Conference on Artificial Intelligence (IJCAI). It was sponsored by sixteen different technical societies (along with some of their subgroups) from the United States, Europe, and Japan. About 600 people attended the conference, and sixty-three papers were presented by authors from nine different countries. The papers were collected in a proceedings volume, which was made available at the conference to all of the attendees.

Because of the success of this first conference, it was decided to hold a second one in London in 1971. During the early years, organization of the conferences was rather informal, decisions about future conferences being made by a core group of some of the leaders of the field who happened to show up at organizing meetings. At the 1971 meeting in London, I left the room for a moment while people were discussing where and when to hold the next conference. When I returned, I was informed that I had been selected to be the “czar” of the next meeting – to be held at Stanford University in 1973. Later, a more formal arrangement was instituted for managing the “International Joint Conferences on Artificial Intelligence,” with a President, a Board of Trustees, and a Secretariat.¹ Since the first meeting, IJCAI conferences are held biennially (on odd-numbered years) with the venue alternating (loosely) between North America and the rest of the world. As at the first conference, proceedings are distributed at the conferences. (Some of these can be obtained from various booksellers, and they are available online at the Digital Library of India Web site, <http://202.41.92.139/>.)

One of the oldest “special interest groups” of the Association for Computing Machinery (ACM) is SIGART (the Special Interest Group for ARTificial intelligence). It began publishing a Newsletter in 1966, which (as the SIGART Web site says) “continued in various incarnations (the SIGART Bulletin, Intelligence Magazine) until 2001.” Today, SIGART supports various conferences and workshops, and it organizes meetings in which AI doctoral students can present their nearly finished work to their peers and to senior researchers for comments and critiques.

As the field began to develop its techniques and methods, graduate courses in artificial intelligence were offered at some universities. Accordingly, some of us

who were teaching these courses thought it would be worthwhile to write or edit books about AI. In 1963, Edward Feigenbaum and Julian Feldman, then assistant professors at UC Berkeley, published a collection of early AI and cognitive science papers that had previously appeared in many different places. The volume was called *Computers and Thought* and was required reading for early students of AI (including me).² As Feigenbaum wrote in the 1995 edition, “Some of the papers are as important today for their fundamental ideas as they were in the late 1950s and early 1960s when they were written. Others are interesting as early milestones of fields that have expanded and changed dramatically.”

In 1965, I published a book about neural-network and statistical methods in pattern recognition.³ That book was followed in 1971 by a book about AI search strategies.⁴ Around the same time, other texts were published by James Slagle⁵ and by Bertram Raphael,⁶ both former Ph.D. students of Marvin Minsky at MIT.

In 1969 Marvin Minsky and Seymour Papert published an influential book in which they proved, among other things, that some versions of Rosenblatt’s perceptrons had important limitations.⁷ Some have claimed that the Minsky–Papert book was the cause of a fading interest in neural-network research, but I doubt this. First, Rosenblatt himself began concentrating on other topics well before 1969,⁸ and the success of heuristic programming methods caused a shift of attention (including my own) away from neural networks during the mid-1960s.

In 1965, Donald Michie at the University of Edinburgh organized the first of several invitation-only “Machine Intelligence” workshops. This first one was held in Edinburgh and was attended by American and European researchers. Attendees gave papers at the workshop, and these were all published in a book edited by N. L. Collins and Donald Michie in 1967. A second workshop was held in September 1966, also at the University of Edinburgh. Subsequent workshops were held annually in Edinburgh through 1971. Thereafter, the workshops were held every few years at various venues. Each workshop resulted in a book with the title *Machine Intelligence N*, where *N* denotes the workshop and volume number.⁹ The last few volumes have been published online by the *Electronic Transactions on Artificial Intelligence*.¹⁰ These books contain some of the most cited and important papers in the early history of the field.

These years saw the United States engaged in war in Vietnam, and Congress wanted to make sure that research supported by the U.S. Defense Department was relevant to military needs. Responding to these pressures, on November 19, 1969, Congress passed the “Mansfield Amendment” to the Defense Procurement Authorization Act of 1970 (Public Law 91-121), which required that the Defense Department restrict its support of basic research to projects “with a direct and apparent relationship to a specific military function or operation.” On March 23, 1972, the Advanced Research Projects Agency was renamed the Defense Research Advanced Projects Agency (DARPA) to reflect its emphasis on projects that contributed to enhanced military capabilities. (The name reverted to the Advanced Research Projects Agency in 1993 and then back to the Defense Advanced Research Projects Agency in 1996.)¹¹

On the other side of the Atlantic, British AI researchers experienced their own funding crisis. One of the U.K.’s main funding bodies for university research, the

Science Research Council, asked Professor James Lighthill, a famous hydrodynamicist at Cambridge University, to undertake an evaluative study of artificial intelligence research. Lighthill's report, titled "Artificial Intelligence: A General Survey," somewhat idiosyncratically divided AI research into three categories, namely, advanced automation, computer-based studies of the central nervous system, and the bridges in between. He called these categories A, C, and B, respectively. Although he came out in favor of continued work in categories A and C, he was quite critical of most AI basic research, including robotics and language processing, which he lumped into category B. He wrote that "In no part of the field have the discoveries made so far produced the major impact that was then [around 1960] promised." He concluded that AI's existing search techniques (which worked on toy problems) would not scale to real problems because they would be stymied by the combinatorial explosion.¹²

Lighthill's report resulted in a substantial curtailment of AI research in the United Kingdom. In particular, one of its casualties was work on FREDDY the robot and other AI work under Donald Michie at Edinburgh. Here is one of Michie's later comments about the effects of the report:¹³

Work of excellence by talented young people was stigmatised as bad science and the experiment killed in mid-trajectory. This destruction of a co-operative human mechanism and of the careful craft of many hands is elsewhere described as a mishap. But to speak plainly, it was an outrage. In some later time, when the values and methods of science have further expanded, and those of adversary politics have contracted, it will be seen as such.

DARPA's shift to shorter term applied research, together with the Lighthill report and criticisms from various onlookers, posed difficulties for basic AI research during the next few years. Nevertheless, counter to Lighthill's assessment, many AI techniques did begin to find application to real problems, launching a period of expansion in AI applications work, as we'll see in the next few chapters.

Notes

1. See <http://www.ijcai.org/IJCAItrustees.php>. [202]
2. Edward A. Feigenbaum and Julian Feldman, *Computers and Thought*, New York: McGraw-Hill Book Co., 1963. (The McGraw-Hill volume is now out of print; it is now available through AAAI Press/MIT Press, 1995 edition.) [203]
3. Nils J. Nilsson, *Learning Machines: Foundations of Trainable Pattern-Classifying Systems*, New York: McGraw-Hill Book Co., 1965; republished as *The Mathematical Foundations of Learning Machines*, San Francisco: Morgan Kaufmann Publishers, 1990. [203]
4. Nils J. Nilsson, *Problem-Solving Methods in Artificial Intelligence*, New York: McGraw-Hill Book Co., 1971. [203]
5. James R. Slagle, *Slagle, James!* his book on AI *Artificial Intelligence: The Heuristic Programming Approach*, New York: McGraw-Hill Book Co., 1971. [203]
6. Bertram Raphael, *The Thinking Computer: Mind Inside Matter*, New York: W. H. Freeman, 1976. [203]
7. Marvin Minsky and Seymour Papert, *Perceptrons: An Introduction to Computational Geometry*, Cambridge, MA: MIT Press, 1969. [203]
8. See Frank Rosenblatt, J. T. Farrow, and S. Rhine, "The Transfer of Learned Behavior from Trained to Untrained Rats by Means of Brain Extracts. I," *Proceedings of the National Academy of Sciences*, Vol. 55, No. 3, pp. 548–555, March 1966. [203]

9. The series maintains a Web page at <http://www.cs.york.ac.uk/mlg/MI/mi.html>. [203]
10. See <http://www.etai.info/mi/>. [203]
11. See http://www.darpa.mil/body/arpa_darpa.html. [203]
12. The text of the report, along with commentary and criticism by leading British AI researchers, was published in 1972 in James Lighthill *et al.* (eds.), *Artificial Intelligence: A Paper Symposium*, London: Science Research Council of Great Britain, 1972. [204]
13. Donald Michie, *Machine Intelligence and Related Topics: An Information Scientist's Weekend Book*, p. 220, New York: Gordon and Breach Science Publishers, 1982. [204]

