SIGART. *Stuart C. Shapiro, Chair, SIGART, Dept. of Computer Science and Center for Cognitive Science, State University of New York at Buffalo, 226 Bell Hall, Buffalo, NY 14260-2000.*

Scope. SIGART, the Special Interest Group for Artificial Intelligence, was formed as a SIC (Special Interest Committee) in 1966 and was converted to a SIG in 1968. It is one of the largest SIGs in ACM and is the only international AI membership organization. SIGART is interested in all aspects of artificial intelligence: theoretical, experimental, and applied; symbolic and nonsymbolic. Its members include students, teachers, researchers, and practitioners.

The *SIGART Newsletter* (first called the *SICART Newsletter*) was published from December 1966 to October 1989. The successor quarterly *SIGART Bulletin* began in January 1990, and has published regularly since. The *Bulletin* contains reviewed papers, correspondence, book and conference reviews, bibliographies, conference announcements and news.

SIGART has nine locals, in the U.S., Canada, and Turkey, plus three "informal" locals and three groups in the process of becoming accredited. SIGART provides its locals some financial support, and is interested in the formation of additional locals. Each local has regular meetings, with local and nonlocal speakers discussing various aspects of AI.

SIGART cosponsors and cooperates with a variety of AI conferences and is willing to increase that number. It also provides direct grants to conferences and workshops, frequently to support student attendance. In 1994, SIGART cosponsored the ACM Symposium on Principles of Database Systems, the Symposium on Lisp and Functional Programming, the Conference on Computational Learning Theory, and the International Conference on Information and Knowledge Management. Conferences held in 1994 in cooperation with SIGART included: the International Workshop on Next-Generation Information Technologies and Systems, the International Conference on Industrial Applications of Artificial Intelligence and Expert Systems, the Golden West International Conference on Intelligent Systems, the International Conference on User Modeling, and the Knowledge-Based Software Engineering Conference.

Surveys and Tutorials. The *SIGART Bulletin* publishes occasional annotated bibliographies and survey articles, and regular topical special issues. Recent annotated bibliographies were on distributed artificial intelligence (Vol. 3, No. 3, 1992) and temporal and modal logic programming (Vol. 5, No. 3, 1994). Recent survey articles were "A review of artificial intelligence in simulation" and "Goals, issues and directions in machine learning of natural language and ontology", both in Vol. 2, No. 1 (1991). Recent special issues were on implemented knowledge representation and reasoning systems (Vol. 2, No. 3, 1991); integrated cognitive architectures (Vol. 2, No. 4, 1991); temporal reasoning (Vol. 4, No. 3, 1993); inductive logic programming (Vol. 5, No. 1, 1994); time in neural networks (Vol. 5, No. 3, 1994); and planning agents (Vol. 6, No. 1, 1995).

Tutorials on AI are regularly given at the International Joint Conference on Artificial Intelligence (IJCAI), which is held in odd-numbered years, and at the U.S. National Conference on Artificial Intelligence, sponsored by AAAI, which is held every year except when IJCAI is in North America. Unfortunately, tutorial materials are generally limited to copies of slides and are usually not available after the tutorial is given. SIGART encourages the submission of written versions of these tutorials to COMPUTING SURVEYS.

Artificial Intelligence.[1] *Artificial Intelligence* (AI) is that field of computer science concerned with the computational understanding of what is commonly called intelligent behavior and with the creation of artifacts that exhibit such behavior. The AI researcher may be characterized as pursuing one or more of the following goals:

Computational psychology. Understanding human intelligent behavior by creating computer programs that behave in the same way people do, using the same algorithms and data structures that the human mind actually uses.

Computational philosophy. Formation of a computational understanding of human-level intelligent behavior, without being restricted to the algorithms and data structures that the human mind actually does (or conceivably might) use. Answering the question, "Is intelligence a computable function?"

Advanced computer science. Pushing outwards the frontier of what we know how to program on computers, especially in the direction of tasks that people can perform, although we don't yet know how to program them. This goal led to one of the oldest definitions of AI: "The attempt to program computers to do what until recently only people could do."

The scope of AI may further be appreciated by briefly considering some of its major subareas:

Natural language. The goal is to form a computational understanding of how people learn and use their native languages, and to produce a computer

[1]This section is a revised and shortened version of S. C. Shapiro's, Artificial Intelligence. In *Encyclopedia of Artificial Intelligence, Second Edition*, S. C. Shapiro, Ed. New York: Wiley, 1992, 54–57. A different and longer revision appeared as S. C. Shapiro, Artificial Intelligence, In *Encyclopedia of Computer Science, Third Edition*, A. Ralston and E. D. Reilly, Eds. Van Nostrand Reinhold, New York, 1993, 87–90.

program that can use a human language at the same level of competence as a native human speaker.

Problem solving and search. Problem solving is concerned with finding or constructing a solution to a problem, where a problem is generally considered to be a space of potential solutions that must be searched to find the correct one or the best one. The AI area of search is concerned with efficient and effective procedures for finding solutions in such spaces.

Knowledge representation and reasoning. Knowledge representation is the area of AI concerned with symbolic languages for representing the information used by intelligent systems. Reasoning is the use of stored explicit knowledge represented in those languages, possibly with some implicit knowledge, to produce additional explicit knowledge.

Learning. Learning is the production of new knowledge other than by deductive inference. AI researchers study all styles of learning: rote learning, conditioning, inductive inference, the design and analysis of experiments, etc.

Vision. Vision, or image understanding, is the interpretation of images that fall on the human retina or the focal plane of a camera. The actual scene being viewed can be two-dimensional, such as a photograph or printed or handwritten text, or three-dimensional, such as the world around us.

Robotics. The area of robotics is concerned with autonomous artifacts that can move about in the actual physical world and/or that can manipulate other objects in the world.

Applications. Throughout the existence of the field, AI research has produced spinoffs into other areas of computer science, from pointer-based data structures, recursive programming, and logic programming, to software development environments and object-oriented programming.

Recently, more and more AI applications are being used by the general public. The fact that they are applications of AI is seldom noticed. What is noticed is an improvement in some service, or that some object (a car, camera, copier, oven, etc.) is becoming more "intelligent."

Some specific AI application successes are cited in the next section. For more examples and more details, see the proceedings of the American Association for Artificial Intelligence (AAAI) Innovative Applications of AI Conferences, and the proceedings of the IEEE Conferences on AI Applications.

Exciting New Developments. This section lists a few of the many successful AI applications.[2]

—"An estimated 70–80 percent of the Fortune 500 companies use AI technology to varying degrees. ... In 1993, the global market for AI systems was

[2]There is only space for a few of the examples contributed by the prominent AI researchers who were solicited for input. No denigration is meant of the examples that were not included. Mention of a specific company, product, or system is for concreteness, and is not to be understood as implying that that particular company, product, or system is necessarily the best in its class. The source of each example is given in brackets.

estimated at about $900 million. ... AI systems proved their strategic value in support of operations Desert Shield and Desert Storm. For example, DART (Dynamic Analysis and Replanning Tool) solved the logistical nightmare of moving the U.S. military assets to the Saudi Desert. ... This one application alone reportedly more than offset all the money the Advanced Research Projects Agency has funneled into AI research in the last 30 years. ... In the data management area ... an AI program automatically processes and indexes news wires into almost 700 categories for Reuters News Service. For this system alone, savings were estimated at more than $1.25 million in a recent year." [*Critical Technology Assessment of the U.S. Artificial Intelligence Sector*, U.S. Dept. of Commerce, Office of Industrial Resource Administration, Strategic Analysis Division, August, 1994.]

—A family of systems based on the CLASSIC representation system is being used by a large number of engineers to configure more than a dozen AT&T telecommunications products; to date, these systems have processed more than two billion dollars' worth of orders. (Brachman)

—The Ricoh "AI Copier" can copy English documents and include in the copy translations into Japanese of difficult English words. It can also take a hand-drawn line drawing and produce a neatened-up version with straight lines, true right angles, etc. (Hart)

—A NASA system manages the scheduling of people and other resources used to process the space shuttles at the Kennedy Space Center, and saves approximately $1M per flight over prior methods. (Friedland)

—A system called Intellipath is routinely used in hundreds of hospitals and medical schools to assist pathologists with diagnosis and surgical decisions. Intellipath uses a large (40,000 nodes) probabilistic network as its knowledge base. (Horvitz via Pearl)

—Over 6,000 copies of Carnegie Group's Testview are being used by companies like Pitney Bowes, Siemens, GE Medical, and Bailey Controls for troubleshooting office, medical, and manufacturing systems. (Fox)