AI Magazine Volume 18 Number 1 (1997) (© AAAI)

A Retrospective of the AAAI Robot Competitions

Pete Bonasso and Thomas Dean

■ This article is the content of an invited talk given by the authors at the Thirteenth National Conference on Artificial Intelligence (AAAI-96). The piece begins with a short history of the competition, then discusses the technical challenges and the political and cultural issues associated with bringing it off every year. We also cover the science and engineering involved with the robot tasks and the educational and commercial aspects of the competition. We finish with a discussion of the community formed by the organizers, participants, and the conference attendees. The original talk made liberal use of video clips and slide photographs; so, we have expanded the text and added photographs to make up for the lack of such media.

There have been five years of robot competitions and exhibitions, with the first event held in San Jose, California, in 1992 and the fifth held in Portland, Oregon, in 1996. Since that first show, we have seen 30 different teams compete, and almost that many more exhibit their robots. The event has become a key attraction of the national and international conferences. In this article, we look back on the form and function of the five years of exhibitions and competitions and attempt to give you some insights into its history, the technical developments, and the political and cultural issues that the organizers and teams faced over the years. We also try to give you a glimpse into the community and camaraderie that develops among the teams during the competitions.

History

We begin with a quick history of the development of intelligent robots to give a better sense of where the competition robots fit in this development. We found it helpful to draw comparisons with developments in aviation. The comparisons allow us to make useful parallels with regard to aspirations, motivations, successes, and failures.

Dreams

Flying has always seemed like a particularly elegant way of getting from one place to another. The dream for would-be aviators is to soar like a bird; for many researchers in AI, the dream is to emulate a human. There are many variations on these dreams, and some resulted in several fanciful manifestations. In the case of aviation, some believed that it should be possible for humans to fly by simply strapping on wings. Early seventeenth-century clockwork automata similarly sought to mimic the superficial aspects of humans, without any real understanding of how the brain or the rest of the body worked.

There are obvious practical motivations for building robots. People have long yearned for mechanical servants to do their instant bidding and yet not burden their consciences. Such servants would do all the dirty work and never complain. They would be humanlike in their ability to understand what their human masters want, superhuman in their ability to carry out their tasks, and yet somehow unaware or joyfully accepting of their lot in life. Karel Capek's play entitled *R.U.R.* (Rossum's Universal Robots) featured humanlike robots that were designed to free humans from drudgery (Asimov 1990). However, the robots didn't like the drudgery any more than the humans; so, they revolted and fought for their freedom. We are still not sure what types of task robots should perform or how they might fit into society.

The Wright Brothers first set the world's record for the longest hang-glider flight, and ROCKY II, a behavior-based robot, was the forerunner of the new Mars Rover Program.

Not all robotics researchers are focused on reproducing human capabilities. Some researchers dream of building mechanical devices that mimic the behavior of simpler biological organisms such as insects. Most people don't see any need for mechanical insects in a world infested with the biting, stinging, disease-carrying sort, but for some, it is a grand challenge to build a device as adaptable and resourceful as the common cockroach.

In the early years of flight, engineers tried to mimic the superficial aspects of winged flight. Ignorant of the subtleties of aerodynamics, the early attempts often failed dramatically. In aviation, the hang glider was the first step as people such as England's Percy Pilcher and Germany's Otto Lilienthal experimented with the lift and thrust forces evident in bird's wings. In robotics, the six-axis PUMA arm epitomized the initial step toward practical robots in industry. However, in both disciplines, there were a myriad of useful and not-so-useful digressions. Early legged robots looked almost as pathetic as the early attempts of humans to fly by flapping strapped-on wings.

Early Successes

Eventually, spurred by requirements in the government sector, early successes were achieved in both arenas. The Wright Brothers first set the world's record for the longest hang-glider flight, and ROCKY II, a behavior-based robot, was the forerunner of the new Mars Rover Program. Driven by military requirements, powered flight soon followed, but it was some time before airplanes were used as a reliable tool in industrial and military applications. Balloons were already being used for reconnaissance, so it was natural for the military to use the more maneuverable airplane for the same purposes. Mail delivery, crop dusting, passenger services, and a host of other applications followed suit, but it took time for the fledgling aircraft industry to prove itself up to the task.

We consider the robots of the competition representative of this early success era, and with funds provided by the Defense Advanced Research Projects Agency (DARPA) and the National Aeronautics and Space Administration (NASA), the caliber of robots increased with each year of the competition with such robots as SRI International's FLAKEY (in competition) and IS Robotics GENGHIS (in exhibition). Just as aviation made strides in its early success era, such as Lindbergh's transatlantic flight, the American Association for Artificial Intelligence competition produced robots with surprising skills (figure 1). CARMEL won the 1992 search-and-navigate event with an as yet unmatched fast sonar-processing algorithm (Congdon et al. 1993). Stanford University's DERVISH won the Office Navigation event in 1994 by maintaining multiple possible states and pruning them with disambiguating sensor data (Nourbakhsh, Powers, and Birchfield 1995). LOLA won the 1995 Trash Pickup event with a seamless integration of navigation and manipulation (Gutierrez-Osuna and Luo 1996). CAIR-2, winner of the 1995 navigation event, would accept navigation directions from humans and ask for and accept help from humans when it got lost, all by means of spoken language (Yang et al. 1996).

However, in contrast to early flight, some of the first steps in building useful robotic tools were easier to achieve. In particular, simple manipulation, painstakingly orchestrated under human control but flawlessly and indefinitely repeated, was achieved with great precision early on by robot welders. Sustained flight required at least a rudimentary understanding of aerodynamics and the invention of powerful yet light power plants. Interestingly, the success of human-powered flight would have to wait until the development of superstrong, superlight space-age materials.

In robotics, one of the main technological stumbling blocks with robots outside the factory, termed *field robots,* was the need for sensing devices and the know-how to interpret the data that they produce. Despite the availability of relatively inexpensive digital cameras, mobile robotics didn't really take off until the introduction of two devices: (1) the sonar transducer, made widely available by the development of automated focusing systems in cameras, and (2) the infrared sensor, used in automatic door openers and a host of other common pieces of equipment.

The sonar and infrared sensors were cheap and plentiful, were relatively easy to interface to computers, and provided a mobile robot with information about the distance to nearby objects. Unfortunately, these devices provided data that were noisy and difficult to interpret. The underlying physics was relatively simple at the coarse level required for applying these devices. The problem was filtering out noise and fusing the information from multiple devices or multiple readings from the same device to provide useful interpretations of the data for navigation and obstacle avoidance. This problem turned out to be unexpectedly complex.

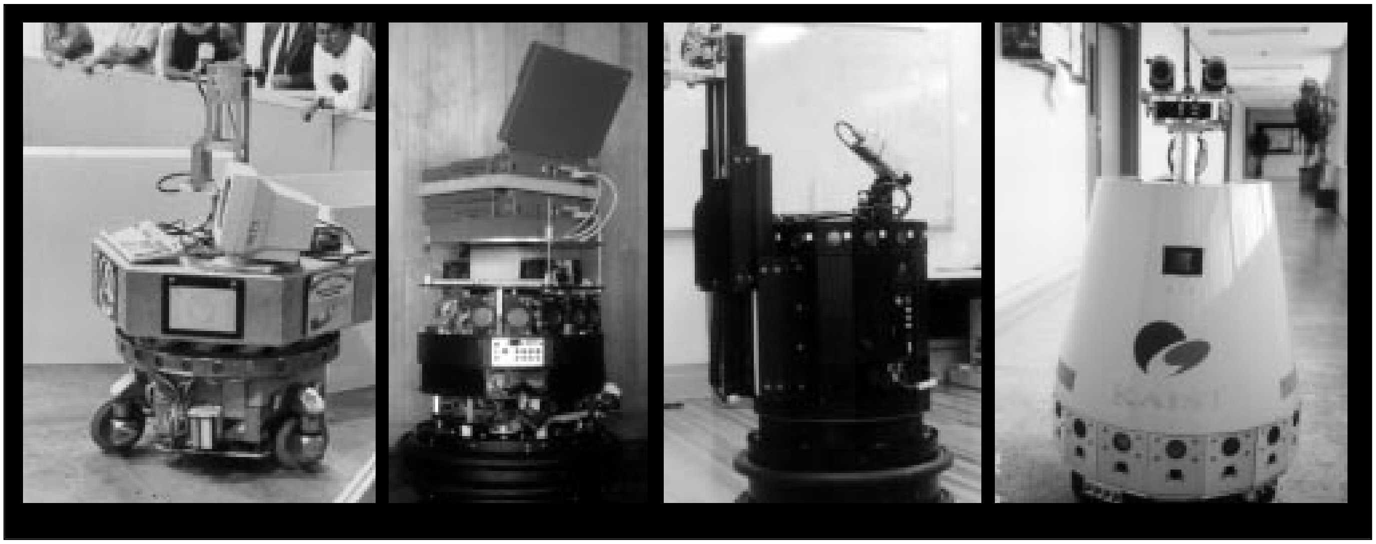

Figure 1. Early Successes in the Robot Competition.

Aviation had the transatlantic flight of Charles Lindbergh. The robot competition had its area and office arenas, which produced winners of navigation events such as the University of Michigan's CARMEL and Stanford University's DERVISH; trash-recycling robots such as North Carolina State University's LOLA; and the smooth-mover CAIR-2, from Korea's AI Institute, which used spoken language.

Real Successes

As we can recall real progress in aviation, we can draw similar parallels with intelligent robots. Real successes in the aviation business included large-capacity passenger jets and faster intercontinental travel. Real success for robots is being realized in the service industry, where such robots as TRC's HELPMATE (see www.helpmaterobotics.com) fetch and deliver meals in semistructured environments such as hospitals. Alongside the high-flying *X15*, we can boast the real-time tracking and grasping of floating objects in zero-g in NASA's EVA Retriever Program (Norsworthy 1994). As humankind moves flight into low-Earth orbit and possibly back to the Moon, ROCKY 4 has gone to Mars, and DANTE (see maas-neotek.arc.nasa.gov/dante/) has explored areas of the Earth too dangerous for human exposure.

Today, we can circumnavigate the globe without landing, fly at incredible speeds, and launch humans and their habitats into space. Human-powered aircraft can span the channel separating England and France. Robotics and aviation have combined to produce aircraft such as the Boeing 777 that can perform many of their operations without human intervention. Such aircraft use sophisticated sensors and control systems. The fact that the air spaces navigated by these systems are relatively free and clear and the landing strips highly engineered makes this sort of automation possible.

Mobile robots are beginning to exhibit interesting behavior in both natural and human-engineered environments. In these cases, however, the opportunities for unanticipated events are much greater than in the commercial airline case. It's this sort of open endedness that makes the realization of the dream of fully autonomous robots so difficult.

Future Dreams

The dreams of both aviators and builders of intelligent robots have been tempered with time and experience, but futurists and science fiction writers keep us wishing for the ultimate in interstellar flight, for example, X-Wing Fighters and starships, and the ultimate in humanoid robots (C3PO), androids (Commander Data), and holographic agents (the Holographic Medical Program of the *Star Trek Voyager* television series).

Power and weight are the factors imposing the greatest limitations on today's aircraft and spacecraft designers. It is unlikely that personal aircraft will soon supplant the family car, nor is it likely that we will be taking a vacation on Mars any time during our lifetimes. Power and weight are limiting factors for robots as well, but advances in sensors, effectors, and the software to make use of these devices are the key to the next generation of robotics and the advent of consumer robotics appliances.

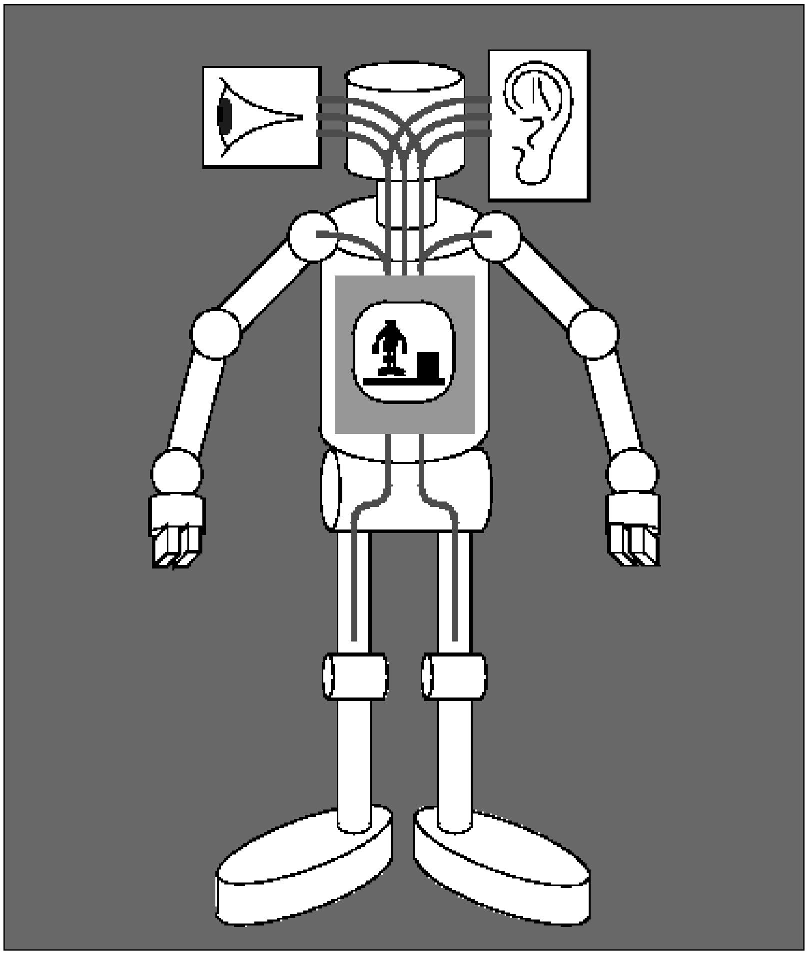

Figure 2. Caricature of an Early Robot.

All the computation—planning, perception, and control—is centralized and relies on complex representations in an attempt to maintain an accurate model of the world. Following the early Asimov novels, we depict the "brains" as being in the robot's belly.

Just as the starships of science fiction are still remote fantasies, the science fiction robots such as Commander Data or even R2D2 are also still a dream. Oddly enough, however, we might see a reasonable facsimile of R2D2 before we see a mechanical insect that comes close to matching the agility, adaptability, and sheer reproductive capacity of everyone's favorite household pest, the cockroach. To a significant extent, the competitions were aimed more toward developing an R2D2 facsimile than a mechanical cockroach, but as it turns out, it is a bit tricky to find an everyday task that humans can do that a properly motivated cockroach cannot.

Technical Challenges

In the following discussion, we highlight some of the technical innovations and scientific contributions that have punctuated the robot competitions over the last five years. We do so using the organization of software and hardware as a theme. Our presentation takes some liberties with historical flow and summarizes a number of complicated issues for the sake of a coherent presentation.

Some of the earliest AI-based robots used centralized computing, with sensors providing data and effectors waiting for input. The amount of data moving from the sensors to the centralized computing resources was often significant. There were architectural considerations concerning software, but these considerations were motivated more by the requirements of high-level deliberation than by the needs of control (figure 2).

Often, these robots relied on building and maintaining complex representations of the environment. These representations were motivated not by the task at hand but by a particular technology concerned with general representations of the world. This motivation turned out to be misguided from the standpoint of tasks tackled in the competition. The disadvantages of general representations were pretty well acknowledged even in 1992 at the first competition, but we still saw a mixed bag of generalists and specialists, and the advantages and disadvantages were illustrated in a fairly dramatic way. Although the movement away from general representations was considered healthy, the resulting degree of specialization was viewed with some alarm.

The computational requirements for maintaining complicated representations were so extreme that competitors often resorted to performing computations off board using radio links to communicate between on-board and off-board computing resources (figure 3). This reliance on off-board computing inevitably came back to haunt the competitors. Existing radio frequency (RF) communication devices turned out to be notoriously error prone because of additional sources of RF signals in the competition environment: cellular phones, security devices, and the equipment used by television crews.

Even ignoring the computational overhead, it was difficult to keep the complex representations in synch with the real world. These baroque representations turned out to be impractical and unnecessary. Not only was the representation difficult if not impossible to maintain, the associated planning systems attempted to use similarly rich representations for planning and prediction. Uncertainty made the underlying dynamics difficult to model and the representations of dubious value for most tasks.

The phrase *look and lurch* characterized the behavior of these earliest attempts, the idea being that the robot gathered a large amount of sensor information; processed it for some time while it stood still; and then made a move of relatively short duration, hoping that the world would not change too much during the move. Needless to say, these robots were not capable of smooth fluid motion.

The next stage of development was characterized by simpler representations tailored to the particular tasks at hand (figure 4). This change in representation resulted in smaller computational requirements. Coupled with improved technology in computers and batteries, on-board computing became practical, and dependence on unreliable RF communications became less widespread. Some amount of RF communication continues to this day, if only to exert emergency control over the robot or maintain a convenient user interface for debugging and demonstration purposes.

This stage also marked the first steps toward distributed computing. In particular, robot manufacturers such as Real World Interfaces (RWI), Denning, and TRC were marketing sensor systems with self-contained processing and standard communication buses. These first attempts at plug-and-play modularity were primarily concerned with handling timing calculations, analog-to-digital conversion, and rudimentary filtering and sensor fusion.

Robot bases, the primary components purchased off the shelf, had always provided rudimentary plug-and-play capability. In the case of bases, however, the local processing supported simple closed-loop control, allowing a remote computer to send high-level commands such as “move forward until the wheels have completed 100 revolutions” over a simple bus such as RS232. In cases where the control loop relied on sensors external to the base, this sort of command-level interface introduced problems resulting from communication delays. More and more competitors were becoming aware of the frequency requirements of accurate and smooth control. These issues were well known in other areas of robotics, but their importance was only realized in AI when trying to combat the performance issues coupled with look-and-lurch-type robots.

As feedback loops were tightened, the look-and-lurch robots gave way to more graceful competitors. Examples of such competitors using commercial advances in the 1992 competition were UNCLE BOB and BUZZ, two Denning robots, and CARMEL, a Cybermotion robot (Dean and Bonasso 1992). In future competitions, these eventually gave way to smaller faster robots such as those made by RWI and Nomadics.

In the next stage (figure 5), the movement toward distributed processing, simpler internal representations, and tighter feedback continued. Many types of behavior were implemented without any intervention from a central computing resource. Central computing became occupied with deliberations requiring more time, such as those involved in building maps, planning paths, and monitoring execution.

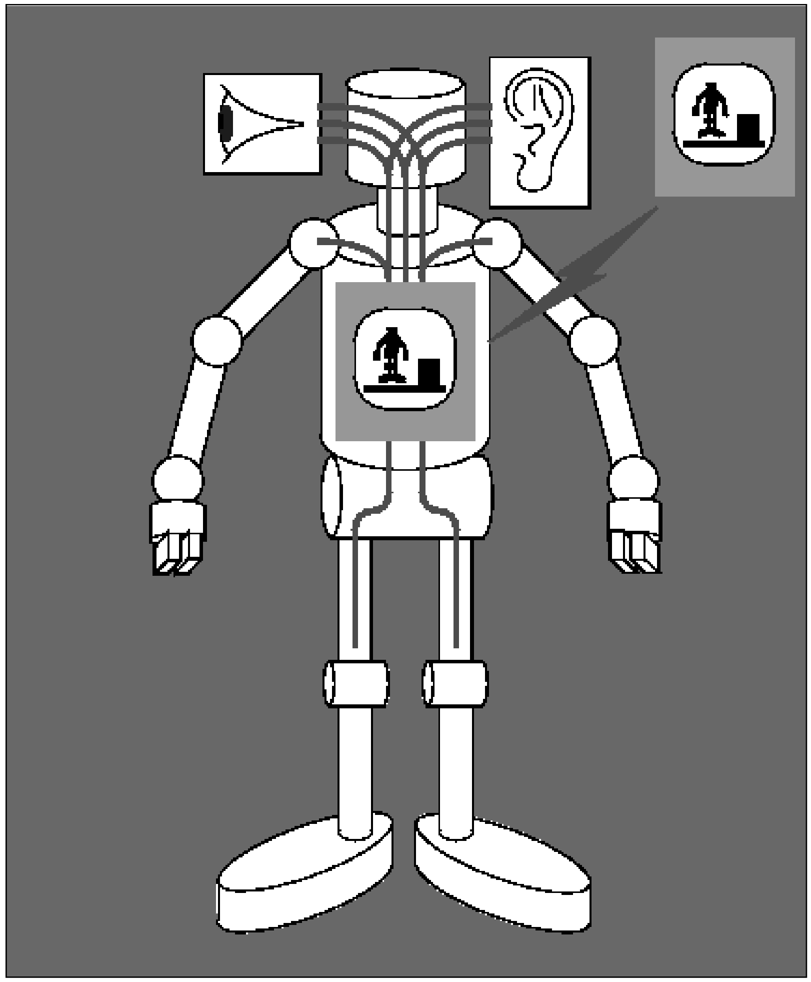

Figure 3. A Variant of the Robot Depicted in Figure 2.

It requires so much computation that some of it has to be carried out on remote computing machinery, with the additional complication of transferring sensor data and commands back and forth over a relatively low-bandwidth radio link.

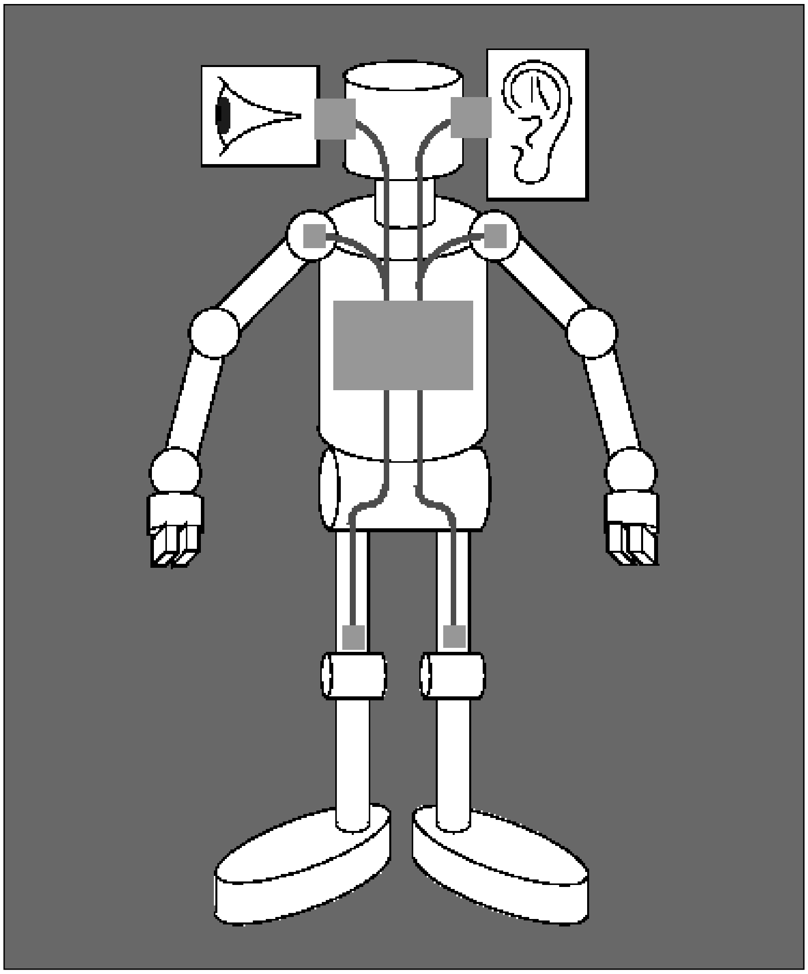

Figure 4. In This Robot, the Computations Have Become Less Centralized and the Representations More Tailored to Particular Tasks.

Sensor data are often interpreted, at least partially, before forwarding to the computer responsible for planning and control.

Some robot builders took this idea to the extreme. In one case, the robot called SCARECROW, built by Dave Miller and his son Jacob, had no computer and instead utilized trivial electromechanical feedback loops to generate relatively simple but remarkably effective behavior given the task at hand (Yeaple 1992) (figure 6).

The community responded to the success of such simple robots by rethinking the tasks used in the competitions. The objective, after all, was to exhibit rudimentary intelligence, whatever that meant. Competition organizers tried to find tasks that were within the capabilities of the existing technology but sophisticated enough not to admit to simple, stateless solutions, that is, solutions that require little in the way of memory, inference, and coordination.

Today, there are many competing approaches, but some general principles have informed the designs. Simple representations are tailored to specific tasks. Layered software allows behaviors such as obstacle avoidance to coordinate smoothly with behaviors such as path following. In addition, there have been some remarkable advances in hardware. Plug-and-play subsystems that combine sensors and effectors are much more common (figure 7).

Manipulators are also much more common. The ones used in the competition are not general-purpose manipulators but devices tailored to a particular set of tasks. The design of manipulators evolved under the pressure of competition tasks meant to require robots to interact with their environment in increasingly complex ways. Early competitions added manipulation tasks but allowed robots to take penalties in lieu of actually performing the required manipulations (termed *virtual manipulation*). Subsequent competitions raised the ante by increasing the penalties, thereby encouraging competitors to build and employ manipulators as well as more sophisticated object recognition and sensing.

The result is that by 1995, we saw some amazingly sophisticated sensing and manipulation. Carnegie Mellon University's (CMU) XAVIER used a V-shaped device designed to connect to the corners of box objects, which could hoist a box overhead, thus allowing unimpeded navigation to a deposit site. CHIP, a small RWI robot from the University of Chicago, integrated sophisticated color vision with a simple gripper taken from a HERO robot. North Carolina State University's LOLA used a commercial manipulator designed specifically for its Nomadics robot, which integrated smoothly as a plug-and-play component. Finally, teams from the University of Minnesota and the Massachusetts Institute of Technology and Newton Labs, inspired by the cockroach syndrome mentioned earlier, used small bug robots with simple reflex claws or collector devices to push or herd objects to the deposit site (figure 8). Some robots, such as CLEMENTINE from the Colorado School of Mines, continue to rely on cheaper, graduate-student labor for their virtual manipulation.

In the early years, there was a lot of talk about architectures and specialized languages. Today, competitors have their preferences, but there are no prospects for a detailed consensus. An architectural tower of Babel eventually gave rise to a set of general principles regarding rough levels of computation based on sampling requirements. In the community of researchers who try to imbue robots with both robustness and intelligence, there is consensus that a layer of software is needed to mediate between the continuous activity of the implemented control laws and the state-based reasoning of the higher functions. In general, however, integrating slow deliberation with fast execution and control requires some art and some science, and the search is on for a better understanding of these issues.

We have characterized this process as evolutionary, but we use artistic license, making the story easier to tell. In fact, each of the architectures represents a family of robots almost like a particular species. It might be that each species has associated with it a *niche problem*, that is, a set of tasks it is particularly well adapted to solve. It might also be that some species are not well adapted for any imaginable set of tasks. The important point is that our technology has adapted to the competition tasks, and the competition tasks have adapted to the demands and expectations of the community of researchers building the robots and observing them in the competitions. The result is an evolving specification for an increasingly capable set of robotic technologies.

We remarked earlier about the tension between existing technological capabilities and the desire to tackle difficult tasks requiring some degree of intelligence. Someone once suggested that we have a competition in which a robot is required to navigate through a crowd of people, pick out the current president of the American Association for Artificial Intelligence (AAAI), and pluck a $100 dollar bill from his or her raised hand. However, this task specification needs considerable refinement and even then might result in unintended consequences, for example, when a flying robot using a sensor that identifies U.S. Treasury ink hacks off the hand of a horrified AAAI president. The specification of common tasks has been both the boon and the bane of robot competitions from the first. How do competition organizers and participants come up with the tasks for the competition? To answer this question, it helps to know the cast of characters involved in the competitions.

Political and Cultural Issues

Consider for a moment some of the different groups that have an interest in the robotics competition. We can roughly divide these groups into three categories: (1) the audience, (2) the backers, and (3) the participants. The audience can be further divided into the conference attendees and the press. AAAI and the International Joint Conference on Artificial Intelligence (IJCAI) are scientific conferences: The attendees expect to be educated about the state of the art; they expect to see significant contributions to science and engineering; and they expect to be entertained, although they might not admit to this. Competitors soon learned that an element of showmanship is important. The conference attendees constitute a sophisticated and critical audience. They want to know what's going on "under the hood," but they also expect the tasks to be well motivated and the performance to be exciting. The prospect of mingling with the robots and robot builders was a big draw for many conference attendees.

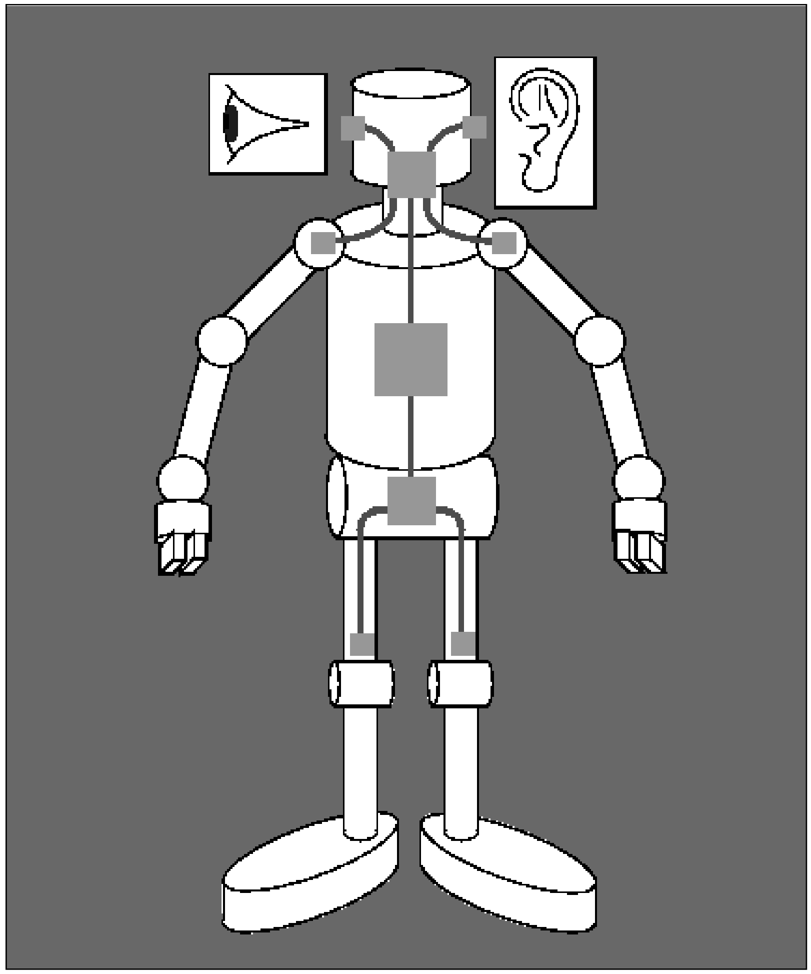

Figure 5. This Robot Design Continues the Progression Started in Figure 4.

This architecture allows for sensor-effector control loops that operate with minimal and only intermittent supervision from a centralized computing resource.

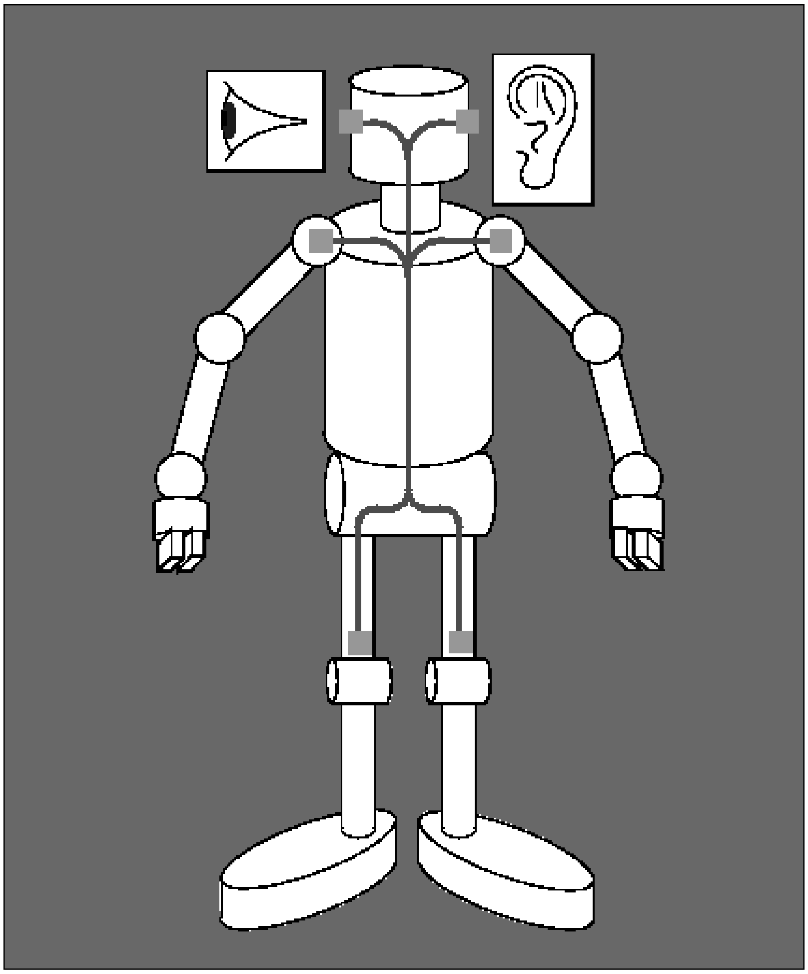

Figure 6. Here is the Extreme Culmination of the Architectural Progression That Began in Figure 4.

This architecture requires no central planning and control processing, and even the computations required for control loops are of the simplest sort, requiring little or no computing by carefully designing the sensors.

For the press, robots are interesting and, at least at some level, understandable to their viewers or readers. Reporters are typically looking for whiz-bang rocket science coupled with competitive drama; technological glitz; human interest; and relevance to their viewers, for example, how robots can make our lives better. However, as a rule, reporters are impatient, skeptical, and generally unwilling to listen to excuses when things don't work on cue. More than once, an organizer has been embarrassed to see how an hour's worth of videotaping and interviews is turned into a three-minute spot on the evening news.

It is important to realize that the competitions are live and online. It's not like giving a demonstration in your lab where you control the environment, have a small army of people to fix code or hardware, and have an audience that is captive and generally patient; you can take a coffee break to distract the visitor while machines are rebooted, bugs fixed, or hardware repaired. Even the research community has been misled by videos that either splice video footage or show only the successful runs and, even then, with dull material excised or, with only slight apologies, the sequence sped up to make it appear more impressive. As an example of this "live" experience, in at least three of the past years, the organizers have tried to show the press the robots that were performing well in the trials, only to have every robot fail in its task.

The backers can be further divided into professional organizations such as AAAI and IJCAI, government agencies and private industries that fund the research, and equipment manufacturers that build robots and sensors and support participants with equipment and technical know-how. AAAI, through its president and executive council, wants to promote the science, increase attendance at the conferences, and obtain good press that ultimately will have a positive impact on Congress and industry.

The government agencies and industrial sponsors expect that those researchers they fund will do well, and they might be on the lookout for groups that currently are not being funded or whose proposals are under consideration. Despite claims to the contrary by conference organizers, the competition can appear to funders like a "bake off" and a convenient way of culling "dead wood" from their programs. We don't know of anyone who has lost funding as a consequence of failure in a competition, but certainly, competitors have touted their success in seeking and justifying funding.

The equipment manufacturers that build robots and sensors want to see their equipment prominently displayed in the winner's circle; their name shown in lights; and their technology applauded and then, of course, purchased in large quantities. We'll have more to say about this area in Commercial Development.

We left the competition and exhibition participants until last because they are in the spotlight, and they are buffeted by the expectations of all the other groups (the organizers of the competition are pretty much in the same boat as the participants and typically compete in those years in which they do not serve as organizers). Competitors want to show off their technology, but they also want to win—it is a competition after all. They want to push the state of the art and, in particular, show progress on basic problems from the previous year. They want the accolades of the press and the attendees, but they also want the scientific appreciation of their peers. Finally, they want to learn and participate in a community effort whose value is enormous and difficult to quantify.

In short, it is the organizers' job (and here again we group the chairs for the competition with the participants because it is a collective enterprise) to create a competition that satisfies all these groups. A large part of this job is specifying a set of tasks for the robots that is both entertaining and has technical merit.

Science and Engineering

How do you push "the envelope" of intelligent robotics? How do you specify a task or set of tasks that is just beyond what we can currently do in the hope that the competitors will be able to rise to the challenge by the time of the competition? If the task is too easy, then it won't be interesting; if the task is too hard, then the competition will likely be embarrassing—either the participants will show up and fail, or they will choose not to participate for fear of embarrassment.

From the outset, organizers have tried to formulate problems that exercise AI technology. However, what is currently considered mainstream AI will typically play a small, although important, role in a successful competitor. For example, we are not particularly interested in better mechanical design, but mechanical design unavoidably will play a role. When you try to provide a precise formulation for an interesting problem, inevitably you introduce loopholes that admit a trivial (or at least AI-free) solution. In the first competition, Miller's SCARECROW demonstrated that a good 90 percent of one of the tasks could be done by a robot without a brain, that is, with only the simplest sort of programming.

Looking back at the first competition, it's surprising how much technology was involved. The tasks were, for the most part, navigation tasks, the last of which required some kind of deliberative reasoning to be competitive. Most of the robots built maps in one phase to use in another, a form of memory-based learning. CARMEL (University of Michigan), ODYSSEUS (CMU), and CHIP (University of Chicago) used machine vision to identify the objects of search in the environment. FLAKEY (SRI), BUZZ (Georgia Institute of Technology), and UNCLE BOB (MITRE Corporation) combined deliberative and reactive methods. Several of the robots used speech to inform the observers about the status of their tasks. In addition, ODYSSEUS used spoken-language recognition. In a look over the years, it didn't seem like we added much in the way of current technologies but, rather, improved on those rudimentary versions displayed in the first competition.

Specifying problems to get to the heart of the problems that AI claims to be interested in is subtle. The effort to produce such specifications made clear to us all the elusive character of intelligence. We wanted competi-

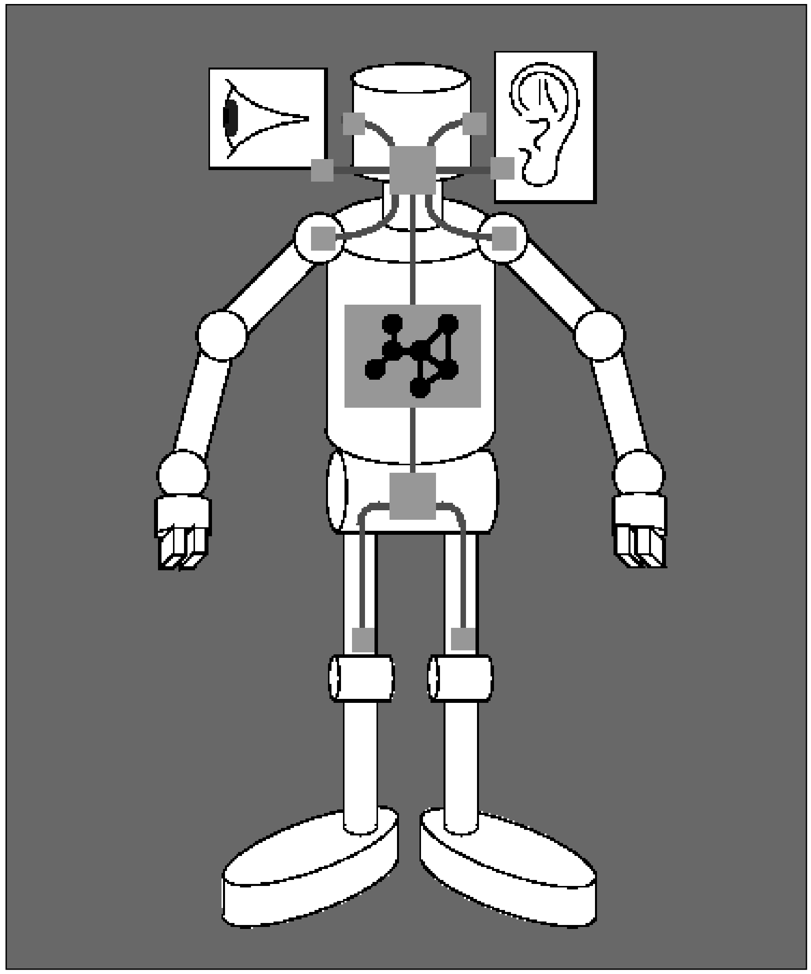

Figure 7. This Robot Depicts Architectural Trade-Offs Established after Five Years of Competitions.

Control and sensing are distributed, with many small computers performing specialized tasks. Planning is centralized, but the representations used in planning are carefully tailored to the robot's tasks. The level of planning detail is limited to the distributed computational units. Uncertainty is dealt with by interacting with the world where appropriate and deliberating about the future where predictions are possible, and the advantages of planning ahead warrant the effort.

Figure 8. Robots That Grab You.

From left to right, Carnegie Mellon University's XAVIER, CHIP from the University of Chicago, and a mobot from the University of Minnesota and the Massachusetts Institute of Technology–Newton Labs.

tors to build systems that used maps and planned and replanned in the face of failed expectations, but these capabilities are, in fact, pretty sophisticated and are not necessary for lots of everyday activities. We discovered, however, that such capabilities were important in tasks in which there was uncertainty, easily accessible regularity, and structure in the environment and in which time is important, although not supercritical. The tasks evolved through the years to reflect these characteristics. You might think that formulating such tasks is easy, but you might think again after subjecting your specification to a group of smart researchers trying to find loopholes, that is to say, competitive advantages, in your lax specification.

An example of the evolution of our task specification was the layout of the arenas. The arena for the 1992 competition was basically a large open area with simple obstacles scattered about. A random walk would take you from one location to any other in a short amount of time, a property that SCARECROW exploited to good end. By contrast, the office environments of future years required more sophisticated mapping and path planning to achieve efficient traversal.

Education

For AAAI-91 in Anaheim, California, program chairs Kathy McKeown and Tom Dean asked John Kender from Columbia University to organize a talk on household robots. In the end, Kender had to choose between attending the conference and being with his wife for the birth of their child. Kender made the right choice, but Dean was stuck with the panel. In fact, it turned out to be a lot of fun and a great success, but the biggest impact was serendipitous. It so happened that Jim Slater and David Zhu from the then-fledgling Nomadics Technology were in Anaheim in the hopes of selling some robots. With the blessing of Pat Hayes, then-president of AAAI, we scheduled a demonstration of the Nomadics robot and invited the press. It was a bit contrived—a bag of groceries was wired to one of the robots—but the response was overwhelming. The press loved it; the attendees loved it; and with the possible exception of the press, viewers were asking the right sort of technological questions. This aspect of "mingling with the conference attendees in an unconstrained manner" continues to be an entertaining, as well as a practical, aspect of each year's competition and exhibition.

It was that afternoon that a group of us got together and started work on the first competition, with the blessing of Pat Hayes and AAAI. From the outset, the event was part competition and part exhibition. Kiosks were set up to show videos and glossy stills of robots from all over the world. Attendees got some idea of the breadth and diversity of robotics research.

Education was key, and a ground swell of interest prompted all sorts of related efforts. David Miller went off to found the KISS Institute for Practical Robotics (KIPR) aimed at education using robots. KIPR used robots partly for motivation and partly as a metaphor for understanding interactivity, behavior, and programming. Other researchers devised educational programs to integrate robots into the classroom for every age group. Robot-building labs became an exciting part of the national conference.

Audience participation gave researchers outside robotics some inkling of how hard the problems facing robotics researchers were. Of course, misconceptions persisted, and dealing with the press continued to be a challenge. Having been brought up on *Star Wars* and reinforced with images of Commander Data on *Star Trek,* the press found it difficult to comprehend just how hard the problems were. The potential for embarrassment was always waiting right around the corner.

Commercial Development

In 1992 when the first competition was held, there were already several well-established manufacturers of mobile robots. Denning and Cybermotion were perhaps the best known to the research community, and their robotic platforms were much in evidence in 1992. Others such as Transitions Research Corporation (now called HelpMate Robotics Inc. after its successful line of hospital robots) were less well known to the research community but were making real progress in fielding robots.

Researchers viewing the 1992 competition were also excited by a number of smaller, affordable robots available from some up-and-coming manufacturers. Nomadics, whose successful demonstration in 1991 made it much easier to make the case for a competition, was out in force in 1992; so, also was another manufacturer of relatively small low-cost robots, RWI. There was often a tension between manufacturers and organizers that more often than not concerned money and publicity; everyone was operating on a shoe string, and there were very few big players from the funding agencies. A bad year could easily spell financial ruin for a small company, and publicity at the competition became an important source of contacts. Miraculously, we all remained on speaking terms.

The funding agencies, despite the organizers' initial fears of bake offs and funding games, were amazingly supportive. The competition simply wouldn't have been possible without the support, financial and otherwise, of Erik Mettala of ARPA, Mel Montemerlo and Peter Friedland of NASA, and Howard Moraff from the National Science Foundation.

Over the years, the organizers, the funders, the leadership of various professional organizations, the conference attendees, and the competition participants became part of a community, a community that is perhaps the most important outgrowth of the robotics competition.

Building Community

Most years, the first day finds a conglomeration of robot teams from different labs and different parts of the world, maybe half of whom did not compete before and maybe half of whom rarely have the same team composition as in years past. As the robots are unpacked and the trials begin, typically several teams discover, to their dismay, that their robot has been damaged in shipment, a computer won't boot, or a key vision board has been short-circuited. Some years, they even have trouble getting to the competition, such as when trying to cross the Canadian border to get to IJCAI-95. The supporting hardware companies and even other teams come to their rescue if they can, despite the fact that they are competitors. For example, the Kansas State University team in the 1995 competition lent North Carolina State its hard drive, and North Carolina State went on to win the Office Cleanup event.

The pressure to have their robots perform causes most teams to spend late nights hacking hardware and software either to recover from catastrophes from the day's events or to make their robots do better in the finals. Every year, we see examples of clever people turning a disaster on one day into a success on the next. From the late nights spent in darkened convention halls, the pressure of the competition, the common interests, and the shared pain and frustration emerges a genuine camaraderie that's difficult to understand unless you've experienced it. First and foremost, the competitors and exhibitors are

Figure 9. This Montage Shows Robot Competition Participants from Previous Years.

interested in realizing the dream of intelligent robots, ad this shared dream makes the hardships worthwhile.

More important than the sharing of equipment and software is the sharing of ideas. Many teams used technologies in later years that they learned from winning teams of previous years. The competition was the impetus for a continuing series of gatherings during the fall and spring symposia focused on the myriad issues of marrying AI and robots. Many of the attendees of these symposia have been past and future participants in the competition. By the end of the competition, something of a temporary family has developed, which is exemplified by the now-traditional family photos taken on the last day (figure 9).

The competition seems to bring out many of the best qualities in science and engineering all over the world: clever ideas, energetic minds, a shared problem to solve, perseverance, a willingness to help one another, good sportsmanship, and a good feeling of accomplishment at the end.

We hope we've been able to convey some idea of what the robotics competitions are all about. Making the event a reality is a complex undertaking—anything worthwhile generally is—but the fruits of this undertaking are many and genuine. It is our hope that everyone in AAAI will take the opportunity to participate in one of the future events either as a spectator or as a participant.

References

Asimov, I. 1990. *Robot Visions*. New York: Visual Publications.

Congdon, C.; Huber, M.; Kortenkamp, D.; Konolige, K.; Myers, K.; and Saffiotti, A. 1993. CARMEL versus FLAKEY: A Comparison of Two Winners. *AI Magazine* 14(1): 49–57.

Dean, T., and Bonasso, R. P. 1993. The 1992 AAAI Robot Competition and Exhibition. *AI Magazine* 14(1): 34–48.

Gutierrez-Osuna, R., and Luo, R. C. 1996. LOLA: Object Manipulation in an Unstructured Environment. *AI Magazine* 17(1): 55–62.

Norsworthy, R. S. 1994. Grasping Objects Autonomously in Simulated KC135 Zero-G. In Proceedings of the AIAA/NASA Conference on Intelligent Robots for Factory, Field, Service, and Space, 757–762. Washington, D.C.: National Aeronautics and Space Administration.

Nourbakhsh, I.; Powers, R.; and Birchfield, S. 1995. DERVISH: An Office-Navigating Robot. *AI Magazine* 16(2): 53–60.

Yang, H. S.; Chung, J.; Ryu, B. S.; and Lee, J. 1996. CAIR-2: Intelligent Mobile Robot for Guidance and Delivery. *AI Magazine* 17(1): 47–54.

Yeaple, J. A. 1992. David Miller's Brainless Robot. *Popular Science* 241(6): 102.

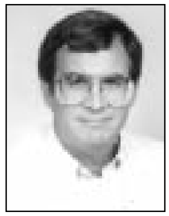

Pete Bonasso is a consulting scientist for AI and robotics at the Automation and Robotics Group of Metrica, Inc. Currently, he is conducting research in the application of intelligent-agent architectures to mobile and manipulator robots in support of the Automation, Robotics, and Simulation Division at the National Aeronautics and Space Administration's Johnson Space Center. He received his B.S. in engineering from the U.S. Military Academy at West Point in 1968 and Master's degrees in operations research and computer utilization from Stanford University in 1974.

Thomas Dean received a B.A. in mathematics from Virginia Polytechnic Institute and State University in 1982 and his Ph.D. in computer science from Yale University in 1986. He is currently professor of computer science and cognitive and linguistic science at Brown University. His general research interests include temporal and spatial reasoning, planning, robotics, learning, and probabilistic inference. His current research is concerned with theories of temporal and spatial inference for reasoning about actions and processes. Of particular interest are problems in which the notion of risk is complicated by the limited time for both deliberation and action. Dean is a fellow of the American Association for Artificial Intelligence; a member of the board of directors of the Computing Research Association; and a member of the board of trustees for the International Joint Conferences on Artificial Intelligence, Inc.