AI Magazine Volume 19 Number 3 (1998) (© AAAI)

Highly Autonomous Systems Workshop

Richard Doyle, Robert Rasmussen, Guy Man, and Keyur Patel

■ Researchers and technology developers from the National Aeronautics and Space Administration (NASA), other government agencies, academia, and industry recently met in Pasadena, California, to take stock of past and current work and future challenges in the application of AI to highly autonomous systems. The meeting was catalyzed by new opportunities in developing autonomous spacecraft for NASA and was in part a celebration of the fictional birth year of the HAL-9000 computer.

In our lifetime, through the eyes of simple robots, grand vistas on other worlds have been unveiled for the first time. Enigmatic questions compel us to go further, to touch these distant landscapes and learn the secrets of the solar system. However, in trying, we find our reach wanting, limited by the link to Earth on which our probes depend. We are learning that to explore further, these probes must go alone, and to go alone, they must become much more intelligent.

The Highly Autonomous Systems Workshop, held in Pasadena, California, on 9 to 10 April 1997, celebrated the birth, in both fact and fiction, of this new generation of explorers. Our goal was to bring together visionaries and skeptics, practitioners and researchers in AI, planetary and space science, spacecraft design, mission design, and mission operations to discuss the important advances in autonomous systems that are propelling this genesis. In exploration, this technology encounters a true and natural test of its maturity where there is no quarter for mediocrity and where all can freely watch and weigh its performance.

Our workshop was one of three symposia in 1997 honoring the fictional birth year of the HAL-9000 computer. The other two were held at the Massachusetts Institute of Technology (MIT) and the University of Illinois at Urbana, which is also the fictional birthplace of HAL. Our meeting was distinguished by reports on actual ongoing work toward building intelligent spacecraft.

Representing government were attendees from the National Aeronautics and Space Administration (NASA), the U.S. Department of Defense (DoD), and the Defense Advanced Research Projects Agency (DARPA).

Our goal was to bring together visionaries and skeptics, practitioners and researchers in AI, planetary and space science, spacecraft design, mission design, and mission operations to discuss the important advances in autonomous systems that are propelling this genesis.

Aerospace attendees included Lockheed Martin, Boeing, and TRW. Academia was represented by institutions such as MIT, Carnegie Mellon University, and the Georgia Institute of Technology.

Featured speakers included one of the founders of AI, Marvin Minsky of MIT, and Louis Friedman, executive director of The Planetary Society. The symposium banquet was graced by two exceptional speakers: the creator of HAL-9000, Arthur C. Clarke, and David Stork of Stanford University, author of *HAL's Legacy*.

The meeting concluded with the announcement of a university design competition on the intriguing topic of *aerobots*, a space platform design that operates in a planetary atmosphere, combining aspects of orbiting platforms and surface vehicles.

The workshop was hosted and organized by the authors and was sponsored by the NASA Autonomy and Information Management Program and the NASA New Millennium Program Autonomy Integrated Product Development Team.

It is our aim, by launching a series of workshops on the topic of highly autonomous systems, to reach out to the larger community interested in technology development for remotely deployed systems, particularly those for exploration. We invite members of this community to join us in helping guide and nurture autonomy's development; learning about its enormous potential, both in space and here at home; sharing ideas; and finding a way to participate.

Overview of the Workshop Sessions

The meeting was organized around four technical sessions: The first session, entitled Historical Visions for Spacecraft Autonomy, took a retrospective view and described a set of visions for spacecraft autonomy that have arisen from different perspectives and evolved over many years of spaceflight. The second session, entitled State-of-the-Art Autonomous Systems, surveyed current autonomous systems work where deployment has already taken place or is well defined and imminent. The third session, entitled Autonomy Technology for 2001, examined autonomous systems research and development for the near term, going out about five years. The fourth and final session, entitled Long-Term Challenges and Benefits of Autonomy, presented bold, unfettered visions of where autonomy technology could reach and what some of its ultimate payoffs might be. A representative sample of the workshop presentations follows.

Mark Brown of the Jet Propulsion Laboratory (JPL) described how deep-space missions have always had drivers for autonomy because of the impracticality of near-continuous communication and the unique difficulties associated with light-time–delayed communication. Examples of long-standing drivers for autonomy on spacecraft include surviving failures, correct execution of time-critical activities (such as achieving orbit), on-board control requiring feedback, and protection of critical resources. A key concept is that the spacecraft must end up in a predictable, commandable state when faults occur. Historically, autonomy has been applied only when deemed necessary, with on-board computing resources being a significant limiting factor.

Bob Connerton of NASA Goddard Space Flight Center spoke about the requirements for autonomy on spacecraft that observe the Earth from orbit. The overwhelming driver is the need to reduce as many as terabits of raw data collected across a diverse set of high-throughput space-based instruments to usable information, sometimes in near real time. On-board feature extraction and data fusion are important capabilities that can support responses to events such as volcanic eruptions or forest fires. Spacecraft will also be arrayed in formations and constellations. These multiple space elements must be controlled precisely and their on-board activities coordinated.

Louis Friedman of The Planetary Society gave a talk entitled "Humans versus Robots, or Where Will the Humans Be?" in which he argued for the value of unmanned spacecraft in performing the basic NASA mission of space exploration. Unmanned spacecraft will always be our first emissaries to remote places, and by making them more autonomous, they can perform new kinds of mission and more in-depth studies and extend our scientific awareness further and further out. In this exciting picture of exploration, a theme is emerging that focuses on the search for life elsewhere—possibly on early Mars, in the suspected subsurface ocean on Jupiter's moon Europa, at planets around nearby stars.

Matthew Barry of the United Space Alliance addressed autonomy from the perspective of NASA's manned spaceflight program. Risk management is the overriding consideration, with human lives being central to the picture. Nonetheless, with cost reduction becoming a major goal, there is considerable interest in capabilities such as decision support to assist flight controllers and astronauts in making procedural choices. In this context, fault diagnosis might actually be less important than sensible reconfiguration decisions, especially those requiring real-time or near–real-time responses. Another area that would benefit greatly from automation support is mission planning and replanning.

Marvin Minsky of MIT gave a talk entitled "What Made HAL Late for His Party?" in which he lamented the lack of definitive progress on questions that the field of AI has long sought to address. He discussed theories of intelligence that address architectural issues, including his "society of mind." He suggested that a theory of which AI techniques actually work, of how well, and in which domains, was probably achievable at this point. He also asserted that a key ingredient for success in NASA's efforts on spacecraft autonomy would be the development of comprehensive models and knowledge bases for on-board use for both engineering and scientific purposes.

Perry McCarty of Lockheed-Martin reported on autonomous control logic for autonomy on underwater submersibles. The basic challenge is to develop an on-board decision-making capacity that can continue a mission in the face of unanticipated, perhaps partially compromising, events. A layered architecture combines a reactive component that monitors conditions and responds to anticipated events with a deliberative component that evaluates vehicle state and capabilities and chooses among courses of action with highest value toward completing the mission, even when the vehicle is found to be in a degraded condition.

Bruce Bullock of ISX spoke on the well-known PILOT'S ASSOCIATE program. The PILOT'S ASSOCIATE is a real-time support system whose job is to efficiently enhance the situational awareness of a pilot in a tactical air battle situation. The task involves modeling the pilot's intentions and state of knowledge and inferring the intents and state of knowledge of friendlies and threats—all while supporting the pilot's actions and communicating information and options accurately and unobtrusively. Plan generation and understanding, information management, and real-time performance are key aspects of this complex human-machine system concept.

Doug Bernard of JPL reported on joint work between NASA Ames Research Center (ARC) and JPL on the REMOTE AGENT, which will be flight tested on the New Millennium Deep Space One mission in 1998. The REMOTE AGENT experiment will demonstrate an autonomy architecture consisting of three reasoning engines and associated models: (1) the *planner-scheduler,* which translates mission goals into a set of on-board activities to be performed, along with the dependencies among them; (2) the *smart executive,* which constructs an explicit timeline of activities and initiates and monitors the execution of these activities; and (3) the *mode identification and reconfiguration system,* which continuously assesses overall spacecraft state, diagnoses faults, and has the authority to command the spacecraft from an incorrect state to the desired state.

Alan Schultz of the Naval Research Laboratory spoke to the core issue of how to test autonomy software, which is of a different nature and complexity from conventional on-board software and will likely require new software-validation concepts. The key idea in this work is to utilize genetic algorithms to explore the space of possible test scenarios, guided by human knowledge of fault classes as a starting point, but avoid subtle biases that can result in inadequate coverage when humans generate the suite of test scenarios. The method has been evaluated for simulated autonomous landing of an F-14 on an aircraft carrier and was able to identify faults not anticipated by the designer.

David Kortencamp of NASA Johnson Space Center described an intriguing application of autonomy to robotic cameras performing inspection tasks of the Space Shuttle or Space Station or in support of astronauts performing extravehicular activities. This autonomy concept is based on a three-tiered architecture whose levels include (1) a *skill manager* for low-level resource management and communication, (2) a *sequencer* for scheduling and monitoring specific activities, and (3) a planner that determines the set of activities to achieve specific goals. Several experiments are planned on the Space Shuttle to validate this technology.

Clark Chapman of the Southwest Research Institute spoke about applications of autonomy for planetary science. Autonomy has a role to play in those situations that involve transient phenomena requiring timely decisions, involve interactive operations in a remote location, and are constrained by limited data rates. On-board autonomy is not appropriate for the highest-level cognitive functions of the scientist but can support data acquisition and data reduction and classification of results in well-defined applications, thereby providing enhanced opportunities where scientists cannot possibly be involved otherwise.

Ron Arkin of Georgia Tech gave an intriguing talk that examined robot design concepts from a suite of unusual viewpoints. Examples included imaginative robots, which simulate and explore the consequences of action before actually performing the action; emotional robots, whose experience of frustration helps them triggers useful mode changes; robots with hormones, which mediate internal communication and control functions; robots that acquire skills using a form of learning analogous to immune system function; and finally, in a rather startling example, a hybrid cockroach-robot system using a grafted microcontroller with potential for applications such as pipe inspection.

Richard Doyle of JPL presented a vision for the development of autonomy technology in which autonomy for science offers strategic value beyond autonomy for engineering or spacecraft functions because science autonomy more directly enables new missions. Scientist-directed on-board software keeps the investigators in intimate contact with the spacecraft, allowing mission priorities to be evolved as scientific understanding of the remote environment evolves, using a combination of conventional algorithms, recognizers to be trained using machine-learning techniques, and knowledge-discovery techniques. Such software is installed at launch time and uploaded during the mission. Ongoing scientist-defined projects in science autonomy include natural satellite search and change detection on planetary surfaces.

Brian Williams of ARC offered a vision for the development of future spacecraft and missions using a model-centric paradigm. The concept starts from the notion of *model-based programming,* where models not only capture knowledge but also are composed directly to realize desired behaviors in the space system. A model-based autonomy kernel can be realized, combining reactive and deductive capabilities, that supports such useful behaviors as anticipation, self-modeling, adaptation, information seeking, and collaboration. Common modeling tools will be essential in realizing this vision, as will new validation techniques, which can themselves draw on model-based concepts.

Gerald Sussman of MIT, in a talk entitled "The Future as I See It," presented a perspective on technology evolution as the development of different kinds of prosthetic, where *prosthetic* is taken in its general sense as a compensator or amplifier for an ability that has been compromised or is inadequate for the task at hand. Different eras have taken different views on what form of machine assistance is most useful. The industrial revolution might be taken as the successful development of mechanical prosthetics. The medical prosthetics emerging today might be just the vanguard of a more general class of biological prosthetics. Sussman explored intriguing concepts for intelligence prosthetics of the future, the logical successor in the sequence of prosthetics evolution.

Clarke left us with the following paraphrase of Descartes, perhaps to be uttered someday by an intelligent machine: "I think, therefore I am, I think."

Banquet Speakers

The workshop was honored with two exceptional banquet speakers, both of whom represented the theme of celebrating 1997 as the fictional birth year of the HAL-9000 computer. The first speaker was none other than the creator of HAL, the author of *2001: A Space Odyssey,* Arthur C. Clarke. Clarke prepared a videotaped greeting to the workshop attendees from his residence in Sri Lanka. The second speaker was David Stork, author of *HAL's Legacy,* which examines the technology prophecies of 2001 from the perspective of the present, offering a number of delightful surprises.

Clarke organized his address to the workshop attendees around a set of reminiscences about HAL and Clarke's own personal interactions with NASA. He recounted how his career has spanned the origins of spaceflight, from the development of rocketry theory, through the realization of spaceborne telecommunications satellites, a concept he first articulated, to the active exploration of the solar system. He lamented the common misinterpretation of the basis of the name *HAL,* reminding us that the correct derivation is heuristically programmed algorithmic computer, one that "has the best of both worlds." He bemused us with how casual, tongue-in-cheek remarks (in this case, regard-

True success, therefore, must be measured in the eagerness of the world to adopt autonomy. However, ironically, the greatest obstacle to this progress is autonomy's own basic nature.

ing the image of a "face" on Mars) are quickly taken up by the less responsible media. He challenged us to explore Jupiter's moon Europa, speculations concerning which have appeared in his recent fiction. Clarke left us with the following paraphrase of Descartes, perhaps to be uttered someday by an intelligent machine: "I think, therefore I am, I think."

Stork gave an engaging talk on the topic of the remarkable prescience regarding computer science–based technologies to be found in both the novel and the movie *2001: A Space Odyssey.* His main motivation for writing *HAL's Legacy* was aesthetics. He wanted to use his knowledge of science and technology to offer a deeper understanding of the film and book, thereby enabling the reader to see them more sensitively, with a more educated eye. Stork reexamined Clarke's and Kubrick's vision and found that in some cases, technology maturity has fallen short of their predictions but in others has surged ahead notably. Computer graphics, supercomputing hardware, and computer reliability have surpassed the vision of *2001.* However, speech recognition, language understanding, common sense, and planning have all fallen far short. Clearly, the general level of AI exhibited by HAL-9000 is not likely to be achieved by the year 2001. However, a computer has now defeated the world chess champion. The conversational interactions between HAL and the human characters Poole and Bowman are loaded with implications about speech generation, speech analysis, and facial expression analysis technology. Stork unearthed from the archives of AT&T what must be the inspiration for some of HAL's discourse: a tape from the 1960s of an early speech generator reciting a verse of the song "Daisy." He also described current work at the MIT Media Lab and elsewhere on inferring the emotional state of a speaker (as an input to semantic analysis of speech) from inflection analysis of the speech signal and visual analysis of facial images. Stork's book contains many additional fascinating examples and insightful analyses.

University Design Competition on Aerobots

Aerobots are a new space platform concept that combines some of the best aspects of orbiter-style missions and surface-style missions. Specifically, an aerobot is designed to exploit the diurnal thermal cycle of a planetary environment by going aloft once a *sol* (a sol is the term assigned to the day cycle in the local planetary environment) and landing once a sol. In this way, an aerobot achieves in part the wide-coverage aspects of an orbiter mission, which can survey an entire planetary surface, along with the in situ exploration aspects of a surface mission, such as those executed by lander and rover combinations, where scientific experiments are conducted in direct interaction with the planetary environment. Although it is possible to predict with some accuracy where an aerobot might land next (with knowledge of prevailing planetary wind patterns, for example), aerobots sample the planetary environment in a stochastic manner, making it nearly impossible to return to a site after leaving. Aerobots are being conceived for exploration wherever planetary atmospheres are present, including Venus, Mars, Jupiter, and Saturn's moon Titan.

Aerobots will require a significant degree of autonomy. Communication will be problematic within the atmospheres at destinations such as Venus, Jupiter, and Titan. If successful scientific missions are to be achieved there, the aerobot platform must be able to grapple with uncertainty again and again and continue to plan and execute the mission while it goes for long periods without ground support. Path planning with a significant random element will be only one of several unique challenges.

A design competition on aerobots was announced at this workshop, targeted at the university community. The intent is to start a cycle where the submissions from the previous design competition are reviewed at each Highly Autonomous Systems Workshop, and a new design competition is announced. Reid Simmons of Carnegie Mellon University is the coordinator for the aerobot design competition.

The Future of Autonomy

In the novel *2001: A Space Odyssey* several decades ago, whether through brilliant foresight or the whimsy of time, Clarke correctly predicted the turning of the millennium as a pivotal moment in the development of highly autonomous systems. He also predicted the momentous impact this development would have on our future—one that would change forever our views of exploration and the bounds of our experience. At the Highly Autonomous Systems Workshop, we gathered not only to celebrate this great act of prescience but also to share our collective experiences and vision for autonomy.

This workshop demonstrated in one presentation after another the broad interest and investment in this technology present today throughout the aerospace, defense, scientific, and exploration communities. It demonstrated that the ideas, computational power, and conviction to make it work are in place. It also demonstrated that advanced autonomy is viewed seriously as a practical answer to real and pressing needs. This unprecedented confluence of need and readiness heralds an era of enormous possibilities.

Highly autonomous systems will greatly extend the safe and efficient exploration of space by enabling probes to hostile and unpredictable places. They will help us understand our own fragile planet from ocean floor to volcanic peak by guiding fleets of explorers and scrutinizing inexhaustible sources of data. They will enhance our national defenses by placing only artificial eyes and ears in harm's way. They will help save lives in space, in the air, and most importantly, on the ground by providing warnings of danger for everything from malfunctioning systems to tsunamis. All these needs are compelling. Our success in addressing them is not of mere academic interest but, rather, serves a vital societal role.

True success, therefore, must be measured in the eagerness of the world to adopt autonomy. However, ironically, the greatest obstacle to this progress is autonomy's own basic nature. The long term vision of *2001* cast intelligent machines not merely as tools but, more significantly, as partners to the human endeavor, capable of deliberate independent action. This desired property of highly autonomous systems is the essence of the word *autonomy,* but it is what skeptics fear most. Independent action is taken as action that is out of control. Moreover, it is often viewed as a usurpment of human volition—an expensive way to do the wrong thing.

The ultimate challenge to highly autonomous systems will therefore be the happy union of control and independence we are able to concoct so that this technology should find an open invitation to wide use. To this end, future workshops will continue to concentrate on this imperative but visionary aspect of autonomous systems and their successful injection into real-world practical applications. We will follow developments from concept to realization in the field to hard lessons learned, and we will chart the purposeful advancement of the technology, providing a forum for objective appraisal.

The future of autonomy is in your hands. We look forward to hearing from all of you at our next workshop. Additional information can be found at ic-www.arc.nasa.gov/ic/Hal9000/.

Richard Doyle is technical section manager of the Information and Computing Technologies Research Section and assistant program manager for the Autonomy Technology Program at JPL. He received a B.A. in mathematics from Boston University and an S.M. and a Ph.D. in computer science from the Massachusetts Institute of Technology Artificial Intelligence Laboratory. He was U.S. program chair for the International Symposium on Artificial Intelligence, Robotics, and Automation for Space, held in Tokyo in July 1997. He gave the invited talk entitled "The Emergence of Spacecraft Autonomy" at the National Conference on Artificial Intelligence in Providence, Rhode Island, in July 1997. His e-mail address is Richard.Doyle@jpl.nasa.gov.

Robert Rasmussen is the mission data system architect for the Advanced Flight Systems Program at JPL. He holds a B.S. and an M.S. in electrical engineering from Iowa State University and a Ph.D. in electrical engineering and mathematics from Iowa State University. He has extensive experience in spacecraft attitude control and computer systems, test and flight operations, and automation and autonomy—particularly in the area of spacecraft fault tolerance. Most recently, he was cognizant engineer for the Attitude and Articulation Control Subsystem on the Cassini mission to Saturn. His e-mail address is Robert.Rasmussen@jpl.nasa.gov.

Guy Man received his bachelor's degree in engineering and mathematics from the University of Redlands. He received an M.S. in engineering and a Ph.D. in mechanical engineering from Stanford University. He is currently responsible for the validation of breakthrough autonomy technologies to drastically reduce mission operations costs and enable new science missions for the twenty-first century in the NASA New Millennium Program. He is one of the three recipients of the 1997 NASA Software of the Year Award. His e-mail address is Guy.Man@jpl.nasa.gov.

Keyur Patel is the technical group supervisor of the Validation Group in the Avionic Systems Section at JPL. He received his B.S. and M.S. degrees from the California Polytechnic Institute at Pomona. He recently chaired the Spacecraft Modeling and Simulation session at the Sixteenth Digital Avionics System Conference. His e-mail address is Keyur.Patel@jpl.nasa.gov.