AI Magazine Volume 13 Number 4 (1992) (© AAAI)

National Aeronautics and Space Administration Workshop on Monitoring and Diagnosis

Richard J. Doyle

■ *The First National Aeronautics and Space Administration (NASA) Workshop on Monitoring and Diagnosis was held in Pasadena, California, from 15 to 17 January 1992. The workshop brought together individuals from NASA centers, academia, and aerospace who have a common interest in AI-based approaches to monitoring and diagnosis technology. The workshop was intended to promote familiarity, discussion, and collaboration among the research, development, and user communities.*

The First National Aeronautics and Space Administration (NASA) Workshop on Monitoring and Diagnosis was held in Pasadena, California, from 15 to 17 January 1992.[1] The workshop was hosted by the Jet Propulsion Laboratory (JPL) and took place at the Ritz-Carlton Huntington Hotel. The meeting was sponsored by NASA's Office of Aeronautics and Space Technology and NASA's Office of Space Systems Development. The members of the program committee were Richard Doyle of JPL, Kathy Abbott of NASA Langley Research Center, Steve Chien of JPL, Ken Forbus of Northwestern University—Institute for the Learning Sciences, Troy Heindel of the Gensym Corporation (formerly of NASA Johnson Space Center [JSC]), Ben Kuipers of the University of Texas at Austin, Ethan Scarl of Boeing Computing Services, and Monte Zweben of NASA Ames Research Center. Together these researchers represented NASA, the AI community, and the aerospace community.

The goal of the workshop was to bring together workers from NASA centers, academia, and aerospace who have a common interest in monitoring and diagnosis technology. The workshop was intended to promote familiarity, discussion, and collaboration among the research, development, and user communities. Attendance was limited to slightly over 60 participants. Roughly 40 percent of the attendees were from NASA; the remaining participants were about equally divided between the academic and aerospace communities. All NASA centers were represented, and there were attendees from eight universities; five private research laboratories; and eight aerospace companies, including one European aerospace company.

The meeting was organized around 16 invited presentations. Topics included ongoing monitoring and diagnosis applications within NASA, examples of technology transfer from university laboratories to real-world development efforts, state-of-the-art research in model-based reasoning (MBR), and an overview of relevant research and applications activities in the European Space Agency (ESA). The presentations were divided roughly equally between research- and applications-oriented projects.

...the presentations were divided ... between research- and applications-oriented projects

In addition to the planned presentations, the agenda also featured an overview of model-based reasoning techniques by Olivier Raiman and Mark Shirley of Xerox Palo Alto Research Center (PARC), breakout sessions in which attendees assembled in small groups and investigated topics of special interest in greater depth, an invited talk by John Muratore of JSC on the most successful application of AI technology within NASA to date, and an open house at JPL.

The breakout sessions were organized around themes that emerged during the first 1-1/2 days of the workshop. These themes were (1) the maturity and applicability of MBR techniques; (2) methods and tools for constructing models; (3) issues relating to dynamic, continuous systems, (4) issues relating to sensors; and (5) technology transfer. This report concentrates on the breakout sessions. For more information on work presented at the meeting, interested readers are invited to obtain the collected abstracts of the workshop.[2]

Although not a stated central theme of the workshop, there was intense interest—and, in some cases, healthy skepticism about—MBR techniques, and one breakout session was devoted to the state and promise of MBR. Muratore pointed out that the knowledge-based systems developed at JSC that are currently supporting space shuttle missions use associative, rather than model-based, diagnostic reasoning techniques. He argued that MBR techniques are not mature enough as a technology, model construction overhead is too expensive, and associative techniques provide adequate capability.

Johan de Kleer of Xerox PARC countered that MBR techniques are not meant as an alternative to other monitoring and diagnosis techniques but as a complement. One of the advantages of a model-based diagnosis approach is that it is possible to detect unforeseen faults, a capability not inherent in any approach based on an implicit or explicit enumeration of fault types. De Kleer agreed that model construction can be difficult but noted that the overhead involved in developing structural, functional, or causal models is not worse than that associated with developing the types of models used in traditional methods such as numeric or discrete-event simulation. He offered the same observation for model validation. The decision to incur a greater cost for increased capability is one that must be made on an effort-by-effort basis. De Kleer also remarked that recent empirical performance analyses and work on reasoning with tolerances for imperfect models suggest that MBR techniques are maturing and are ready to be applied.

Jane Malin of JSC led the breakout session on model building. MBR practitioners and would-be MBR practitioners recognize that the overhead associated with developing models can be prohibitive. Moreover, models are rarely shared because of either incompatibilities of representation or of intended use. The session helped to sharpen the different viewpoints and wish lists of researchers and users in this area.

Users would like standardized libraries of component models for their application domains. Ideally, such libraries would be rich enough so that system models could be constructed merely by specifying connections among instantiations of various primitive component types. The resulting models should be useful in systems performing application tasks such as simulation, fault detection, or diagnosis.

Barring such an ideal model-building facility, users would like to have standardized modeling languages that are easy to use, provide multiple representational schemes (for example, qualitative, quantitative, Boolean), and provide for domain-specific extensions. They would like to be able to reuse and modify models constructed by others for different purposes. Researchers also want standardized models and model development environments so that research results can be shared and evaluated more easily. They would also like a test suite of standard examples to be established for various purposes, for example, to test the empirical performance of diagnosis algorithms.

The breakout session on coping with issues associated with continuous dynamic systems was led by David Throop of Boeing Computer Services. Among the issues examined were requirements for model accuracy and approaches to model validation. NASA's concept of the orbital replaceable unit (ORU) provides a basis for model accuracy: A model should be accurate enough to support diagnosis resolution to the ORU level. A pragmatic NASA approach to model validation is equally straightforward: Enumerate and test your diagnosis algorithm against a set of known faults established during design analysis. Beyond this level of validation, MBR techniques can diagnose some unforeseen faults; human troubleshooting expertise fills the remaining gaps.

The RTDS effort has been the most successful application of AI technology within NASA to date

Dan Dvorak of AT&T and the University of Texas at Austin described his work on monitoring continuous dynamic systems. He has developed methods for detecting discrepancies in continuous-valued data by comparing sensor readings to the predictions of a qualitative-quantitative simulation model. Some faults are detected directly as discrepant readings, but others are detected as mutually inconsistent readings (even though individual readings are not discrepant). A key idea is that the qualitative-quantitative model enables reasoning with partial quantitative knowledge and reveals all the behaviors consistent with this partial knowledge, thus eliminating some sources of false positives in fault detection.

Another output of this session was a list of possible dissertation topics mapping out open problem areas. Among these topics were a theoretical basis for combining model-based and associative reasoning systems, diagnosis techniques for dynamic systems, noise-filtering techniques, more methods for reducing spurious behaviors in qualitative simulation, and fast or anytime algorithms for diagnosis.

The session led by Ethan Scarl of Boeing Computer Services concentrated on issues relating to sensors. Several workshop participants reported on complementary approaches to *sensor placement*, the problem of choosing a sensor configuration during system design that enables nominal and faulted states to be distinguished in the operational system while minimizing sensor costs such as weight, power consumption, and polling requirements. Scarl, Steve Chien of JPL, and Janos Sztipanovits of Vanderbilt University addressed this criterion for diagnosis. Scarl's method allows a minimal set of sensors to be found for a given resolution of diagnosis. Chien and Sztipanovits also focused on measures of how well and how quickly different sensor configurations distinguish different states. They argued that accuracy and timeliness are as important as distinguishability.

Sensor placement based on "diagnosability" necessarily assumes a set of a priori known faults. NASA missions, however, have exhibited examples of unforeseen faults manifesting well after a space platform has been launched. Doyle and Chien argued that the "monitorability" criteria that are designed to detect anomalies, rather than specific faults, complement the "diagnosability" criteria for sensor placement and help to ensure that sensor data are available to enable the detection of unknown faults, as a precursor to troubleshooting.

Doyle, Chien, and Usama Fayyad of JPL reported on related work on sensor selection for real-time monitoring. They described a set of sensor importance measures used to focus the attention of experienced mission operators interpreting sensor data in real time. Sensor selection is intended to avoid information overload on human operators and to support efficient anomaly detection.

The final breakout session, led by Gregg Swietek of NASA Headquarters, was on technology transfer. NASA applications are unique, and they exhibit demanding capability and performance requirements. Workshop participants heard about a number of these applications. In the workshop's invited talk, Muratore described the history of the real-time data systems (RTDSs) project at JSC, which has delivered a family of knowledge-based systems for space shuttle monitoring and diagnosis applications to the JSC Mission Control Center. The application domains for the RTDS systems range from ground communications to engine performance to wind conditions for shuttle landings. Muratore was himself a flight controller and was in a unique position to champion the work as both a user and a developer. He spoke of two discontinuous events in the technology-transfer process: flash point and freeze point. *Flash point* occurs when the interest of the user community is fired. The users themselves see more and more possibilities, and the developer suddenly faces many more demands. Once flash point passes, users become increasingly comfortable with the new capabilities. Eventually, they reach a

AAAI-93 Robot Exhibition

Preliminary Call for Participation

Following the highly successful robotics exhibition at AAAI-92, AAAI is planning to hold a robot competition at the national conference in Washington D.C. in July of 1993. The purpose of this Call is to advise potential participants of the event, and to solicit input on the format of the exhibition and rules of the competition.Last year's competition was a three-stage event in which mobile robots demonstrated skills of reactivity, exploration, and directed search (a detailed description is in the Summer 1992 issue of *AI Magazine*).

Mobile robotics is an area where much of the research in diverse AI areas can be effectively and creatively combined to give interesting results. At AAAI-93, we would like to extend the competition to highlight as wide a range of robotic research as possible, and to stress the "intelligent" aspects of their behavior. In addition to mobile robots, we are also considering having a competition among robotic manipulators, either stationary or attached to mobile platforms.

If you are interested in participating, and would like to receive more detailed information about the competition, please contact:

Kurt Konolige, Artificial Intelligence Center
SRI International, 333 Ravenswood Avenue
Menlo Park, CA 94025, (konolige@ai.sri.com)
or

Reid Simmons, School of Computer Science
Carnegie Mellon University, 5000 Forbes Avenue
Pittsburgh, PA 15213, (reids@cs.cmu.edu)

point where they would not want to do their job the old way. This event is *freeze point*. The new approach officially becomes part of the accepted way of doing things. System reliability becomes a critical issue; further development ceases; and steps are taken to document, engineer, and officially deliver the new technology.

The RTDS effort has been the most successful application of AI technology within NASA to date and provides one example of successful technology transfer. Several other speakers also described NASA applications or ongoing examples of technology transfer. Chuck Pepe of McDonnell-Douglas Corporation described continuing work on the knowledge-based autonomous test engineer (KATE) system at NASA Kennedy Space Center (KSC). KATE, originally conceived and implemented by Scarl and John Jamieson of KSC, is one of the earliest model-based diagnosis applications. Scott Karro, Janet Lauritsen, and Dennis Lawler of JSC described fault management applications for the Space Station Control Center at JSC. Jean-Michel Darroy of Matra Marconi Space (Toulouse, France) gave an overview of monitoring and diagnosis research and applications projects within ESA, which were remarkable for their similarity to NASA projects.

At the open house at JPL, Eric Biefeld of JPL demonstrated the operations mission planner (OMP) automated scheduling system, which utilizes an iterative refinement technique to support both generative and reactive scheduling. Chien demonstrated the JPL selective monitoring (SELMON) system, which focuses operator attention in support of real-time monitoring by ordering and selecting sensor data. Two additional monitoring and diagnosis applications at JPL were featured: Rick Martin demonstrated SHARP (spacecraft health automated reasoning prototype), and Ursula Schwuttke demonstrated MARVEL (multimission automation for real-time verification of engineering link). Both systems were originally prototyped for the Voyager mission and have since been applied to other JPL missions. Finally, Lawler of JSC demonstrated the digraph-based fault-dependency visualization and analysis tool called FEAT (failure environment analysis tool).

The workshop resulted in members of different communities meeting each other halfway and taking away new ideas and new information. Customers within NASA and the aerospace community learned about the state of the art in AI-based monitoring and diagnosis: the techniques being developed and their levels of maturity as technologies. AI technology developers within NASA and the aerospace community learned about the techniques they can inherit from the research community and about the problems and requirements of NASA user groups. In addition, AI researchers in academia working in the areas of model-based, causal, qualitative, associative, and probabilistic reasoning learned about NASA application domains and the availability of system models and test data.

A Second NASA Workshop on Monitoring and Diagnosis is planned. The second meeting is conceived as more of a working meeting, organized primarily around breakout sessions, with researchers, developers, and users matching approaches to requirements, brainstorming, and discussing issues in great depth.

Notes

1. The meeting described in this report was hosted by the Jet Propulsion Laboratory, California Institute of Technology, and was supported under a contract with the National Aeronautics and Space Administration.

2. The collected abstracts of the workshop can be obtained by writing to Richard Doyle, Jet Propulsion Laboratory, MS 525-3660, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099.

Richard J. Doyle is technical group supervisor for artificial intelligence and assistant program manager for knowledge systems at the Jet Propulsion Laboratory, California Institute of Technology, in Pasadena, California. He received a B.A. in mathematics, with a minor in astronomy, from Boston University in 1980 and S.M. and Ph.D. degrees in computer science from the Massachusetts Institute of Technology in 1984 and 1988, respectively. His research interests lie in model-based reasoning and machine learning. His current work concerns sensor-selection methods for attention focusing and anomaly detection to support operators performing real-time system monitoring.