

Editorial Introduction

Space Applications of Artificial Intelligence

Steve Chien, Robert Morris

■ *We are pleased to introduce the space application issue articles in this issue of AI Magazine. The exploration of space is a testament to human curiosity and the desire to understand the universe that we inhabit. As many space agencies around the world design and deploy missions, it is apparent that there is a need for intelligent, exploring systems that can make decisions on their own in remote, potentially hostile environments. At the same time, the monetary cost of operating missions, combined with the growing complexity of the instruments and vehicles being deployed, make it apparent that substantial improvements can be made by the judicious use of automation in mission operations.*

The case for increasing the level of autonomy and automation for space exploration is well known. Stringent communications constraints are present, including limited communication windows, long communication latencies, and limited bandwidth. Additionally, limited access and availability of operators, limited crew availability, system complexity, and many other factors often preclude direct human oversight of many functions. In fact, it can be said that almost all spacecraft require some level of autonomy, if only as a backup when communications with humans are not available or fail for some reason.

Increasing the levels of autonomy and automation using techniques from artificial intelligence allows for a wider variety of space missions and also frees humans to focus on tasks for which they are better suited. In some cases autonomy and automation are critical to the success of the mission. For example, deep space exploration may require more autonomy in the spacecraft, as communication with ground operators is sufficiently infrequent to preclude continuous human monitoring for potentially hazardous situations.

Space applications of AI can also be divided in terms of three kinds of operations they support: predictable, unpredictable, and real time (Jónsson et al. 2007). Even predictable operations can be extremely complex — enabling artificial intelligence to play a key role in automation to manage complexity or to assist human decision making. Unpredictability of the operating environment increases requirements on the AI system to appropriately respond in a wide range of situations. Real-time requirements may impose limitations on the amount of reasoning performed by the AI system.

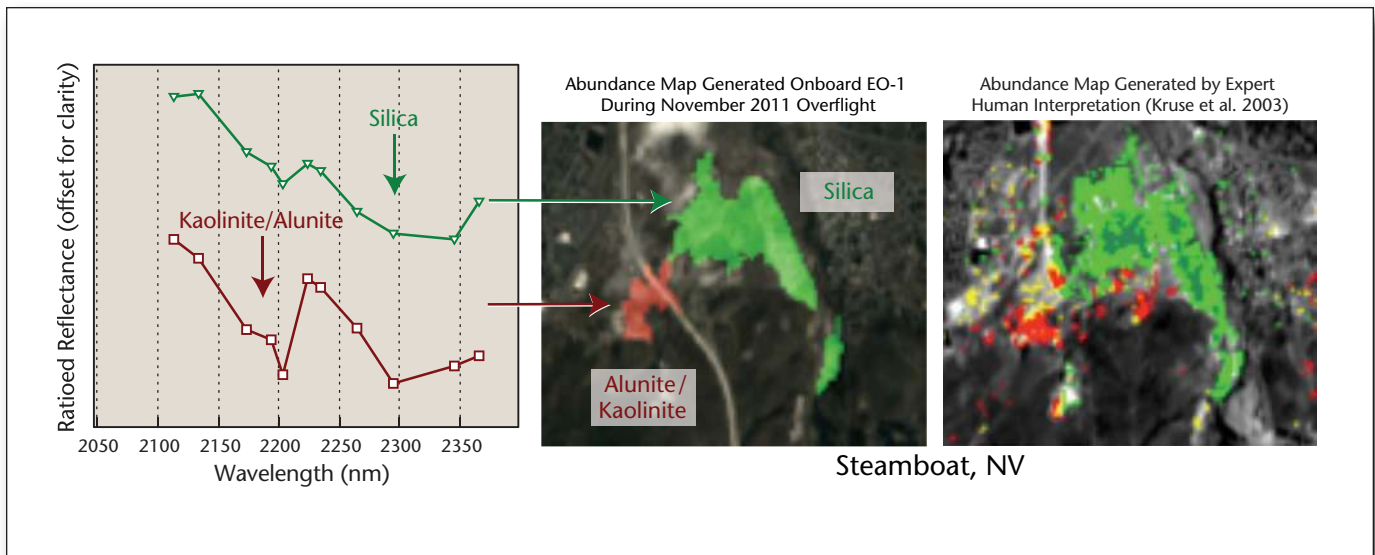


Figure 1. Onboard Spectral Analysis of Imaging Spectroscopy Data During 2011–2012 Demonstrated on EO-1.

Onboard software performed spectral endmember detection and mapping, enabling automatic abundance mapping to reduce data volume by orders of magnitude (Thompson et al 2013). These onboard automatically derived compositional maps (at left) were consistent with prior expert human interpretations (at right).

Of course, deploying large numbers of integrated intelligent agents, each utilizing multiple AI technologies, is the end vision of space AI technologists. The first major step toward this vision was the remote agent experiment (RAX) (Muscettola et al. 1998, Bernard et al. 2000). RAX controlled the Deep Space 1 spacecraft for two periods totaling approximately 48 hours in 1999. RAX included three AI technologies: a deliberative, batch planner-scheduler, a robust task executive, and a model-based mode identification and reconfiguration system.

More recently, the autonomous sciencecraft has controlled the Earth Observing-1 mission for almost 10 years as this article goes to press. This run of operations includes more than 50,000 images acquired and hundreds of thousands of operations goals. The autonomous sciencecraft (Chien et al. 2005a) includes three types of AI technologies: a model-based, deliberative, continuous planner-scheduler (Tran et al. 2004, Rabideau et al. 2004), a robust task executive, and onboard instrument data interpretation including support vector machine-learning derived classifiers (Castano et al. 2006, Mandrake et al. 2012) and sophisticated instrument data analysis (see figure 1) (Thompson et al. 2013b).

Many individual AI technologies have also found their way into operational use. Flight operations such as science observation activities, navigation, and communication must be planned well in advance. AI-based automated planning has found a natural role to manage these highly constrained, complex operations. Early successes in this area include the ground processing scheduling system (Deale et al.

1994) for NASA space shuttle refurbishment and the SPIKE system used to schedule Hubble Space Telescope operations (Johnston and Miller 1994). SPIKE enabled a 30 percent increase in observation utilization (Johnston et al. 1993) for Hubble, a major impact for a multibillion dollar mission. Also impressive is that SPIKE or components of SPIKE have been or are being used for the FUSE, Chandra, Subaru, and Spitzer missions. AI-based planning and scheduling are also in use on the European Space Agency's Mars express and other missions. For a more complete survey of automated planning and scheduling for space missions see Chien et al. (2012a).

In this issue, the article by Mark D. Johnston, Daniel Tran, Belinda Arroyo, Sugi Sorensen, Peter Tay, Butch Carruth, Adam Coffman, and Mike Wallace describes the deep space network (DSN) scheduling engine (DSE) component of a new scheduling system that provides core automation functionality for scheduling of NASA's deep space network, supporting scores of missions with hundreds of tracks every week. The article by Russell Knight, Caroline Chouinard, Grailing Jones, and Daniel Tran describes the application and adaptation of the ASPEN (automated scheduling and planning environment) framework for operations of the Orbital Express (OE) mission.

Because space missions produce enormous petascale data sets, machine learning, data analysis, and event recognition for science and engineering purposes has been another fertile area for application of AI techniques to space applications (Fayyad, Hausler, and Stolorz 1996). An early success was the use

of the decision tree machine-learning techniques in SkiCat (Fayyad, Weir, and Djorgovski 1993) to semi-automatically classify the second Mount Palomar Sky Survey, enabling classification of an order of magnitude greater sky objects than by manual means. Another early advance was the use of Bayesian clustering in the AutoClassAutoClass system (Cheeseman and Stutz 1996) to classify infrared astronomical satellite (IRAS) data. From these beginnings has emerged a plethora of subsequent applications including automatic classification and detection of features of interest in earth (Mazzoni et al. 2007a, 2007b) and planetary (Burl et al. 1998, Wagstaff et al. 2012) remote-sensing imagery. More recently, these techniques are also being applied to radio science signal interpretation (Thompson et al. 2013a).

In this issue the article by José Martínez Heras and Alessandro Donati studies the problem of telemetry monitoring and describes a system for anomaly detection that has been deployed on several European Space Agency (ESA) missions.

Surface missions, such as Mars Pathfinder, Mars Exploration Rovers (MER), and the Mars Science Laboratory (MSL), also present a unique opportunity and challenge for AI. The MER mission uses several AI-related systems: The MAPGEN (Ai-Chang et al. 2004, Bresina et al. 2005) constraint-based planning system for tactical activity planning, the WATCH (Castano et al. 2008) system (used operationally to search for dust devil activity and to summarize information on clouds on Mars.), and the AEGIS system (Estlin et al. 2012) (used for end-of-sol targeted remote sensing to enhance MER science).

Many rover operations, such as long- and short-range traverse on a remote surface; sensing; approaching an object of interest to place tools in contact with it; drilling, coring, sampling, assembling structures in space, are characterized by a high degree of uncertainty resulting from the interaction with an environment that is at best only partially known. These factors present unique challenges to AI systems.

In this issue, the article by David Wettergreen, Greydon Foil, Michael Furlong, and David R. Thompson addresses the use of onboard rover autonomy to improve the quality of the science data returned through better sample selection, data validation, and data reduction.

Another challenge for autonomy is to scale up to multiple assets. While in an Earth-observing context multiple satellites are already autonomously coordinated to track volcanoes, wildfires, and flooding (Chien et al. 2005b, Chien et al. 2012b), these systems are carefully engineered and coordinate assets in rigid, predefined patterns. In contrast, in this issue, the article by Logan Yliniemi, Adrian K. Agogino, and Kagan Tumer tackles the problem of multirobot coordination for surface exploration through the use of coordinated reinforcement learning: rather than

being programmed what to do, the rovers iteratively learn through trial and error to take actions that lead to high overall system return.

The significant role of AI in space is documented in three long-standing technical meetings focused on the use of AI in space. The oldest, the International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS) covers both AI and robotics. I-SAIRAS occurs roughly every other year since 1990 and alternates among Asia, North America, and Europe¹ with 12 meetings to date. Second, the International Workshop on Planning and Scheduling for Space occurs roughly every other year with the first meeting² in 1997 with eight workshops thus far. Finally, the IJCAI³ workshop on AI and space has occurred at each IJCAI conference beginning in 2007 with four workshops to date.

We hope that readers will find this introduction and special issue an intriguing sample of the incredible diversity of AI problems presented by space exploration. The broad spectrum of AI techniques, including but not limited to machine learning and data mining, automated planning and scheduling, multi-objective optimization, and multiagent, present tremendously fertile ground for both AI researchers and practitioners.

Acknowledgements

Portions of this work were carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Notes

1. See robotics.estec.esa.int/i-SAIRAS.
2. See robotics.estec.esa.int/IWPSS.
3. See ijcai.org.

References

- Ai-Chang, M.; Bresina, J.; Charest, L.; Chase, A.; Hsu, J. C.-J.; Jonsson, Ari; Kanefsky, B.; Morris, P.; Rajan, K.; Yglesias, J.; Chafin, B. G.; Dias, W. C.; Maldague, P. F. 2004. Mapgen: Mixed Initiative Planning and Scheduling for the Mars Exploration Rover Mission. *IEEE Intelligent Systems* 19(1): 8–12.
- Bresina, J. L.; Jónsson, Ari K.; Morris, P. H.; and Rajan, K. 2005. Activity Planning for the Mars Exploration Rovers. In *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling*, 40–50. Menlo Park, CA: AAAI Press.
- Bernard, D.; Dorais, G. A.; Gamble, E.; Kanefsky, B.; Kurien, J.; Millar, W.; Muscettola, N.; Nayak, P.; Rouquette, N.; Rajan, K.; Smith, B.; Taylor, W.; and Tung, Y.-W. 2000. Remote Agent Experiment: Final Report. NASA Technical Report 20000116204. Moffett Field, CA: NASA Ames Research Center.
- Burl, M. C.; Asker, L.; Smyth, P.; Fayyad, U.; Perona, P.; Crumpler, L.; and Aubele, J. 1998 Learning to Recognize Volcanoes on Venus. *Machine Learning* 30(2–3): 165–194. [dx.doi.org/10.1023/A:1007400206189](https://doi.org/10.1023/A:1007400206189)

- Castano, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic Detection of Dust Devils and Clouds on Mars. *Machine Vision and Applications* 19(5–6): 467–482. dx.doi.org/10.1007/s00138-007-0081-3
- Castano, R.; Mazzoni, D.; Tang, N.; Doggett, T.; Chien, S.; Greeley, R.; Cichy, B.; Davies, A. 2006. Onboard Classifiers for Science Event Detection on a Remote Sensing Spacecraft. In *Proceedings of the 12th ACM International Conference on Knowledge Discovery and Data Mining (KDD 2006)*. New York: Association for Computing Machinery. dx.doi.org/10.1145/1150402.1150519
- Cheeseman, C., and Stutz, J. 1996. Bayesian Classification (AUTOCLASS): Theory and Results. In *Advances in Knowledge Discovery and Data Mining*, ed. U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy. Menlo Park, CA: AAAI Press.
- Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davies, A.; Mandl, D.; Frye, S.; Trout, B.; Shulman, S.; Boyer, D. 2005a. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication* 2(4): 191–216. dx.doi.org/10.2514/1.12923
- Chien, S.; Cichy, B.; Davies, A.; Tran, D.; Rabideau, G.; Castano, R.; Sherwood, R.; Mandl, D.; Frye, S.; Shulman, S.; Jones, J.; Grosvenor, S. 2005b. An Autonomous Earth-Observing Sensorweb. *IEEE Intelligent Systems* 20(3): 16–24. dx.doi.org/10.1109/MIS.2005.40
- Chien, S.; Johnston, M.; Policella, N.; Frank, J.; Lenzen, C.; Giuliano, M.; Kavelaars, A. 2012a. A Generalized Timeline Representation, Services, and Interface for Automating Space Mission Operations. Paper presented at the 12th International Conference on Space Operations (SpaceOps 2012), Stockholm, Sweden, 11–15 June.
- Chien, S.; McLaren, D.; Doubleday, J.; Tran, D.; Tanpipat, V.; Chitadron, R.; Boonyaroonnet, S.; Thanapakpawin, B.; and Mandl, D. 2012b. Automated Space-Based Monitoring of Flooding in Thailand. Paper presented at the 11th International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS 2012), Turin, Italy, September.
- Deale, M.; Yvanovich, M.; Schnitzuius, D.; Kautz, D.; Carpenter, M.; Zweben, M.; Davis, G.; and Daun, B. 1994. The Space Shuttle Ground Processing Scheduling System. In *Intelligent Scheduling*, ed. M. Zweben and M. Fox, 423–449. San Francisco: Morgan Kaufmann Publishers.
- Estlin, T. A.; Bornstein, B. J.; Gaines, D. M.; Anderson, R. C.; Thompson, D. R.; Burl, M.; Castaño, R.; and Judd, M. 2012. AEGIS Automated Science Targeting for the MER Opportunity Rover. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3): Article 50. doi 10.1145/2168752.216876
- Fayyad, U.; Haussler, D.; and Stolorz, P. 1996. Mining Scientific Data. *Communications of the ACM* 39(11): 51–57. dx.doi.org/10.1145/240455.240471
- Fayyad, U. M.; Weir, N.; and Djorgovski, S. G. 1993. SKICAT: A Machine Learning System for Automated Cataloging of Large Scale Sky Surveys. In *Proceedings of the Tenth International Conference on Machine Learning*. San Francisco: Morgan Kaufmann Publishers.
- Johnston, M.; Henry, R.; Gerb, A.; Giuliano, M.; Ross, B.; Sanidas, N.; Wissler, S.; and Mainard, J. 1993. Improving the Observing Efficiency of the Hubble Space Telescope. In *Proceedings of the 9th Computing in Aerospace Conference*. Reston, VA: American Institute of Aeronautics and Astronautics. dx.doi.org/10.2514/6.1993-4630
- Johnston, M. D., and Miller, G. 1994. Spike: Intelligent Scheduling of Hubble Space Telescope Observations. In *Intelligent Scheduling*, ed. M. Fox and M. Zweben, 391–422. San Francisco: Morgan-Kaufmann Publishers.
- Jónsson, A.; Morris, R. A.; and Pedersen, L. 2007. Autonomy in Space: Current Capabilities and Future Challenge. *AI Magazine* 28(4): 27–42.
- Kruse, F. A.; Boardman, J. W.; and Huntginton, J. F. 2003. Final Report: Evaluation and Geologic Validation of EO-1 Hyperion. Technical Report, Analytical Imaging and Geophysics LLC, Boulder, CO.
- Mandrake, L.; Umaa, R.; Wagstaff, K. L.; Thompson, D.; Chien, S.; Tran, D.; Pappalardo, R. T.; Gleeson, D.; and Castaño, R. 2012. Surface Sulfur Detection via Remote Sensing and Onboard Classification. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3(4): 77. dx.doi.org/10.1145/2337542.2337562
- Mazzoni, D.; Garay, M. J.; Davies, R.; and Nelson, D. 2007a. An Operational MISR Pixel Classifier Using Support Vector Machines. *Remote Sensing of Environment* 107(1): 149–158. dx.doi.org/10.1016/j.rse.2006.06.021
- Mazzoni, D.; Logan, J. A.; Diner, D.; Kahn, R.; Tong, L.; and Li, Q. 2007b. A Data-Mining Approach to Associating MISR Smoke Plume Heights with MODIS Fire Measurements. *Remote Sensing of Environment* 107(1): 138–148. dx.doi.org/10.1016/j.rse.2006.08.014
- Muscettola, N.; Nicola, M.; Nayak, P. P.; Pell, B.; and Williams, B. C. 1998. Remote Agent: To Boldly Go Where No AI System Has Gone Before. *Artificial Intelligence* 103(1) (1998): 5–47.
- Rabideau, G.; Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Mandl, D.; Frye, S.; Shulman, S.; Bote, R.; Szwaczkowski, J.; Boyer, D.; and Van Gaasbeck, J. 2004. Mission Operations with Autonomy: A Preliminary Report for Earth Observing-1. Paper presented at the International Workshop on Planning and Scheduling for Space (IWPSS 2004), Darmstadt, Germany, 23–25 June.
- Thompson, D. R.; Briskin, W. F.; Deller, A. T.; Majid, W. A.; Burke-Spolaor, S.; Tingay, S. J.; Wagstaff, K. L.; and Wayth, R. B. 2013a. Real Time Adaptive Event Detection in Astronomical Data Streams: Lessons from the Very Long Baseline Array. *IEEE Intelligent Systems* 29(1): 48–55. 10.1109/MIS.2013.10
- Thompson, D. R.; Bornstein, B.; Chien, S.; Schaffner, S.; Tran, D.; Bue, B.; Castano, R.; Gleeson, D.; Noell, A. 2013b. Autonomous Spectral Discovery and Mapping onboard Onboard the *EO-1* Spacecraft. *IEEE Transactions on Geoscience and Remote Sensing* 51(6): 3567–3579. dx.doi.org/10.1109/TGRS.2012.2226040
- Tran, D.; Chien, S.; Rabideau, G.; Cichy, B. 2004. Flight Software Issues in Onboard Automated Planning: Lessons Learned on *EO-1*. Paper presented at the International Workshop on Planning and Scheduling for Space (IWPSS 2004), Darmstadt, Germany, 23–25 June.
- Wagstaff, K. L.; Panetta, J.; Ansar, A.; Greeley, R.; Hoffer, M. P.; Bunte, M.; and Schorghofer, N. 2012. Dynamic Landmarking for Surface Feature Identification and Change Detection. *ACM Transactions on Intelligent Systems and Technology* 3(3), article number 49.

Steve Chien is head of the Artificial Intelligence Group and a senior research scientist at the Jet Propulsion Laboratory, California Institute of Technology where he leads efforts in automated planning and scheduling for space exploration. He is a four time honoree in the NASA Software of the year competition.

Robert Morris is a senior researcher in computer science in the Exploration Technology Directorate, Intelligent Systems Division at NASA Ames Research Center. Over the years he has been involved in projects applying AI-based planning and scheduling technologies to solving problems in observation planning, navigation planning and autonomous operations of exploration vehicles.