

A Recap of the AAAI / IAAI 2018 Conferences and the EAAI Symposium

*Sheila McIlraith, Kilian Weinberger, Karen Myers,
Eric Eaton, Michael Wollowski, G. Michael Youngblood*

■ *The 2018 AAAI Conference on Artificial Intelligence, the 2018 Innovative Applications of Artificial Intelligence, and the 2018 Symposium on Educational Advances in Artificial Intelligence were held February 2–7, 2018, at the Hilton New Orleans Riverside, New Orleans, Louisiana. This report, based on the prefaces contained in the AAAI-18 proceedings and program, summarizes the events of the conference.*

The 32nd AAAI Conference on Artificial Intelligence (AAAI-18) was held February 2–7, 2018, at the Hilton New Orleans Riverside, New Orleans, Louisiana. The program chairs were Sheila McIlraith (University of Toronto, Canada) and Kilian Weinberger (Cornell University, USA). Held as an independent program within the AAAI conference was the 30th Innovative Applications of Artificial Intelligence (IAAI-18), February 4–6, and the Eighth Symposium on Educational Advances in Artificial Intelligence (EAAI-18), February 3–4. IAAI-18 was chaired by G. Michael Youngblood (PARC, a Xerox Company) and cochaired by Karen Myers



AAAI-18 Cochair Sheila McIlraith.

(SRI International). The EAAI-18 symposium was cochaired by Eric Eaton (University of Pennsylvania, USA) and Michael Wollowski (Rose-Hulman Institute of Technology, USA).

The purpose of the AAAI conference is to promote research in artificial intelligence and to encourage scientific exchange among AI researchers, practitioners, scientists, and engineers in affiliated disciplines. AAAI-18 had a diverse technical track, student abstracts, poster sessions, invited speakers, tutorials, workshops, and exhibit and competition programs, all selected according to the highest reviewing standards. AAAI-18 welcomed submissions on mainstream AI topics, as well as novel crosscutting work in related areas. IAAI-18 and EAAI-18 were independent programs held under the AAAI-18 umbrella.

AAAI-18

The surge in public interest in AI technologies that we've seen over the past few years continued in 2017–2018 with stories of AI R&D initiatives filling the newswires, and with the societal and economic

impact of AI a point of public and government discussion worldwide. AAAI-18 saw a similar surge in interest, with submissions and attendance numbers breaking the records set in 2017. AAAI-18 continued its tradition of attracting top-quality papers from the breadth of AI. In 2018, we were excited to see increases in submissions across all areas and especially in computer vision and natural language processing, mirroring the worldwide increase in R&D activities related to machine learning and deep learning.

The AAAI-18 program consisted of a core technical program of original research presentations. It additionally featured a broad range of tutorials, workshops, invited talks and panels, student abstracts, What's Hot talks from specialized AI conferences, a debate, and presentations by senior members. The program was rounded out by technical demonstrations, exhibits, the AI job fair, the student outreach program, and a game night. AAAI-18 continued its tradition of colocating with the long-running IAAI conference and the EAAI symposium.

This year AAAI-18 featured the first Oxford-style debate, addressing the controversial statement



AAAI-18 Cochair Killian Weinberger.

“Advances in machine learning have displaced the need for logic in AI.” This somewhat lively debate featured Thomas Dietterich and Bart Selman arguing in favor, and Gary Marcus and Francesca Rossi arguing in opposition. Kevin Leyton-Brown served as moderator with wit and good humor. AAAI-18 also highlighted the emerging topic of human-AI collaboration with a dedicated sequence of technical sessions bookended by 30-minute invited talks. Invited speakers in this emerging topic included Kobi Gal (Ben-Gurion University of the Negev, Israel), Ross Knepper (Cornell University, USA), Devi Parikh (Georgia Institute of Technology, USA), and Matthew Taylor (Washington State University, USA). This initiative was cochaired by Ece Kamar and Julie Shah.

Presidential Address

AAAI-18 featured an impressive lineup of invited speakers. The program began with AAAI President

Rao Kambhampati’s Presidential Address, “Challenges of Human-Aware AI Systems.”

Research in AI, Kambhampati noted, suffers from a long-standing ambivalence to humans — swinging, as it does, between their replacement and their augmentation. Now, as AI technologies enter our everyday lives at an ever-increasing pace, he noted that there is a greater need for AI systems to work synergistically with humans. To do this effectively, he argued, AI systems must pay more attention to aspects of intelligence that help humans work with each other — including emotional and social intelligence.

Kambhampati discussed the research challenges in designing such human-aware AI systems, including modeling the mental states of humans in the loop, recognizing their desires and intentions, providing proactive support, exhibiting explicable behavior, giving cogent explanations on demand, and engen-



Former AAI President Barbara Grosz at AAAI-18.

dering trust. He surveyed the progress made so far on these challenges, highlighted some promising directions, and touched on the additional ethical quandaries that such systems pose. Kambhampati concluded his presidential address by arguing that the quest for human-aware AI systems broadens the scope of AI enterprise, necessitates and facilitates true interdisciplinary collaborations, and can go a long way towards increasing public acceptance of AI technologies.

Invited Talks

Yejin Choi concluded the first day of the conference with her invited talk, “From Naive Physics to Connotation: Learning and Reasoning about the World Using Language.” Intelligent communication, she noted, requires reading between the lines, which in turn requires rich background knowledge about how

the world works. However, she pointed out, learning unspoken commonsense knowledge from language is nontrivial, as people rarely state the obvious, such as “my house is bigger than me.” Choi discussed how we can recover the trivial everyday knowledge just from language without an embodied agent. A key insight is this — the implicit knowledge people share and assume systematically influences the way people use language, which provides indirect clues to reason about the world. For example, from the statement “Jen entered her house,” we can infer that her house is bigger than her.

Choi explained how we can organize various aspects of commonsense — ranging from naive physics knowledge to more abstract connotations — by adapting representations of frame semantics. She then discussed neural network approaches that complement the frame-centric approaches, noted the



AAAI President Subbaro Kambhampati Delivers His Presidential Address at AAI-19.

challenges in current models and formalisms, and pointed to avenues for future research.

On Monday, February 5, Charles Isbell opened the day with his talk, “How Machines Learn Best from Humans.” Isbell noted that “we build machine-learning systems because we want them to behave a certain way.” In this case, he said, the “we” is usually human beings. “Whether we want to convey particular strategies or subtle preferences that define the objective itself,” he observed, “some form of knowledge transfer from person to algorithm is always needed.” Interactive machine learning, he said, “focuses on techniques for facilitating that transfer in the context of solving artificial intelligence problems with machine-learning techniques.” Isbell surveyed some of the problems and techniques studied in interactive machine learning with a special

emphasis on counterintuitive design principles that have arisen from the results of experiments with human participants, especially where those counterintuitive principles arise from *we* being wrong about *us*. A surprise guest appearance by Michael Littman made for a lively and engaging talk.

The evening joint AAI-IAAI talk was presented by Zoubin Ghahramani, the chief scientist at Uber. Ghahramani reviewed the foundations of probabilistic AI, highlighted some current areas of research at the frontiers, and touched on topics such as Bayesian deep learning, probabilistic programming, Bayesian optimization, and AI for data science.

Tuesday’s opening invited talk, “Fair Questions,” was delivered by Cynthia Dwork (Harvard / Radcliffe Institute for Advanced Study). The unfairness of algorithms, she noted — for tasks ranging from advertis-



Invited Talk Audience at AAAI-18.

ing to recidivism prediction — has attracted considerable attention in the popular press. Dwork devoted her talk to the (relatively) new study of a mathematically rigorous theory of fairness: definitions, methods, and provable limits and trade-offs, providing a lens for hot-button policy issues such as interpretability and raising new questions for future research.

Joe Halpern followed later in the day with his talk, “Actual Causality: A Survey.” The problem of defining actual causation, Halpern pointed out, goes beyond mere philosophical speculation. For example, in many legal arguments, actual causation is precisely what needs to be established in order to determine responsibility. The philosophy literature has been struggling with the problem of defining causality since the days of Hume, in the 1700s. Many of the

definitions have been couched in terms of counterfactuals. In 2001, Halpern and Judea Pearl introduced a new definition of actual cause, using Pearl’s notion of structural equations to model counterfactuals. The definition has been revised twice since then, extended to deal with notions like responsibility and blame, and applied in databases and program verification. In his talk, Halpern surveyed the last 15 years of work, including joint work with Judea Pearl, Hana Chockler, and Chris Hitchcock.

The conference’s final invited talk, “How Should We Evaluate Machine Learning for AI?,” was presented by Percy Liang (Stanford University, USA) on Wednesday morning. Liang noted that while machine learning has undoubtedly been hugely successful in driving progress in AI, it implicitly brings with it the train-test evaluation paradigm. This stan-



The 2018 AAAI/EAAI Outstanding Educator Award Is Presented by Rao Kambhampati to Todd Neller.

standard evaluation only encourages behavior that is good on average. It does not ensure robustness as demonstrated by adversarial examples, and it breaks down for tasks such as dialogue that is interactive or that does not have a correct answer. Liang described alternative evaluation paradigms with a focus on natural language understanding tasks, and discussed ramifications for guiding progress in AI in meaningful directions.

Technical Program

The AAAI-18 proceedings includes 938 original research publications, selected by rigorous double-blind review from a pool of 3,800 well-formed original submissions — an all-time record for AAAI. Submissions heralded from all over the world with approximately 48 percent coming from Asia, 29 percent from North America, and 19 percent from

Europe. Of particular note was the 58 percent increase in submissions from China, surpassing the total number of submissions from the US for the first time. The final program reflected an overall acceptance rate of 24.7 percent. Selected papers were presented at the conference in either an oral or a poster format. Technical sessions comprised 13-minute oral presentations and 2-minute spotlight advertisements of poster presentations. The posters were presented during dedicated evening sessions. The assignment of presentation format was based on perceived breadth of interest and topical balance, and did not imply a quality rating.

Four papers presented at AAAI-18 exemplified high standards in technical contribution and exposition by regular and student authors. They were singled out for special mention. The AAAI-18 Outstanding Paper award was presented to Chenjun Xiao,



AAAI Past President Thomas Dietterich Presents AAAI Honors and Special Awards at AAAI-18.

Jincheng Mei, and Martin Müller, for their paper “Memory-Augmented Monte Carlo Tree Search.” An honorable mention was presented to Juan D. Correa, Jin Tian, and Elias Bareinboim for “Generalized Adjustment under Confounding and Selection Biases.”

The AAAI-18 Outstanding Student Paper award was presented to Jakob N. Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson for “Counterfactual Multi-Agent Policy Gradients.” An honorable mention was awarded to Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P. Dickerson, and Vincent Conitzer for their paper “Adapting a Kidney Exchange Algorithm to Align with Human Values.”

Other conference programs included the Senior Member track, cochaired by David Aha and Prasad Tadepalli; the Doctoral Consortium Abstracts pro-

gram, cochaired by Maria Chang and Dan Magazeni; the Student Abstracts program, chaired by Erez Karpas and Nir Lipovetzsky; and the Technical Demonstrations program, chaired by Shirin Sohrabi.

Mounting a conference the size and scope of AAAI is no small feat, requiring the dedicated efforts of a large number of volunteers. Over 2,415 individuals served on the program committee. We also depended on the expertise and commitment of 220 senior program committee members, 56 area chairs, 4 special topic chairs, and 2 emerging topic chairs. We wanted to extend our gratitude to the AAAI-18 conference committee chair, Shlomo Zilberstein, and last year’s program committee cochaairs Satinder Singh and Shaul Markovitch for their sage advice and assistance. The AAAI organization handled all conference arrangements with grace and the highest degree of professionalism.



IAAI Chair Michael Youngblood Welcomes Attendees to IAAI.

Thanks to the volunteers, staff, sponsors, presenters, and attendees, all of whom contributed to making AAAI-18 in New Orleans a success. We look forward to reuniting next year amidst the sun and surf of Hawaii — to share new ideas, engage in discussion and debate, and through our collective efforts, to continue to advance the field of artificial intelligence.

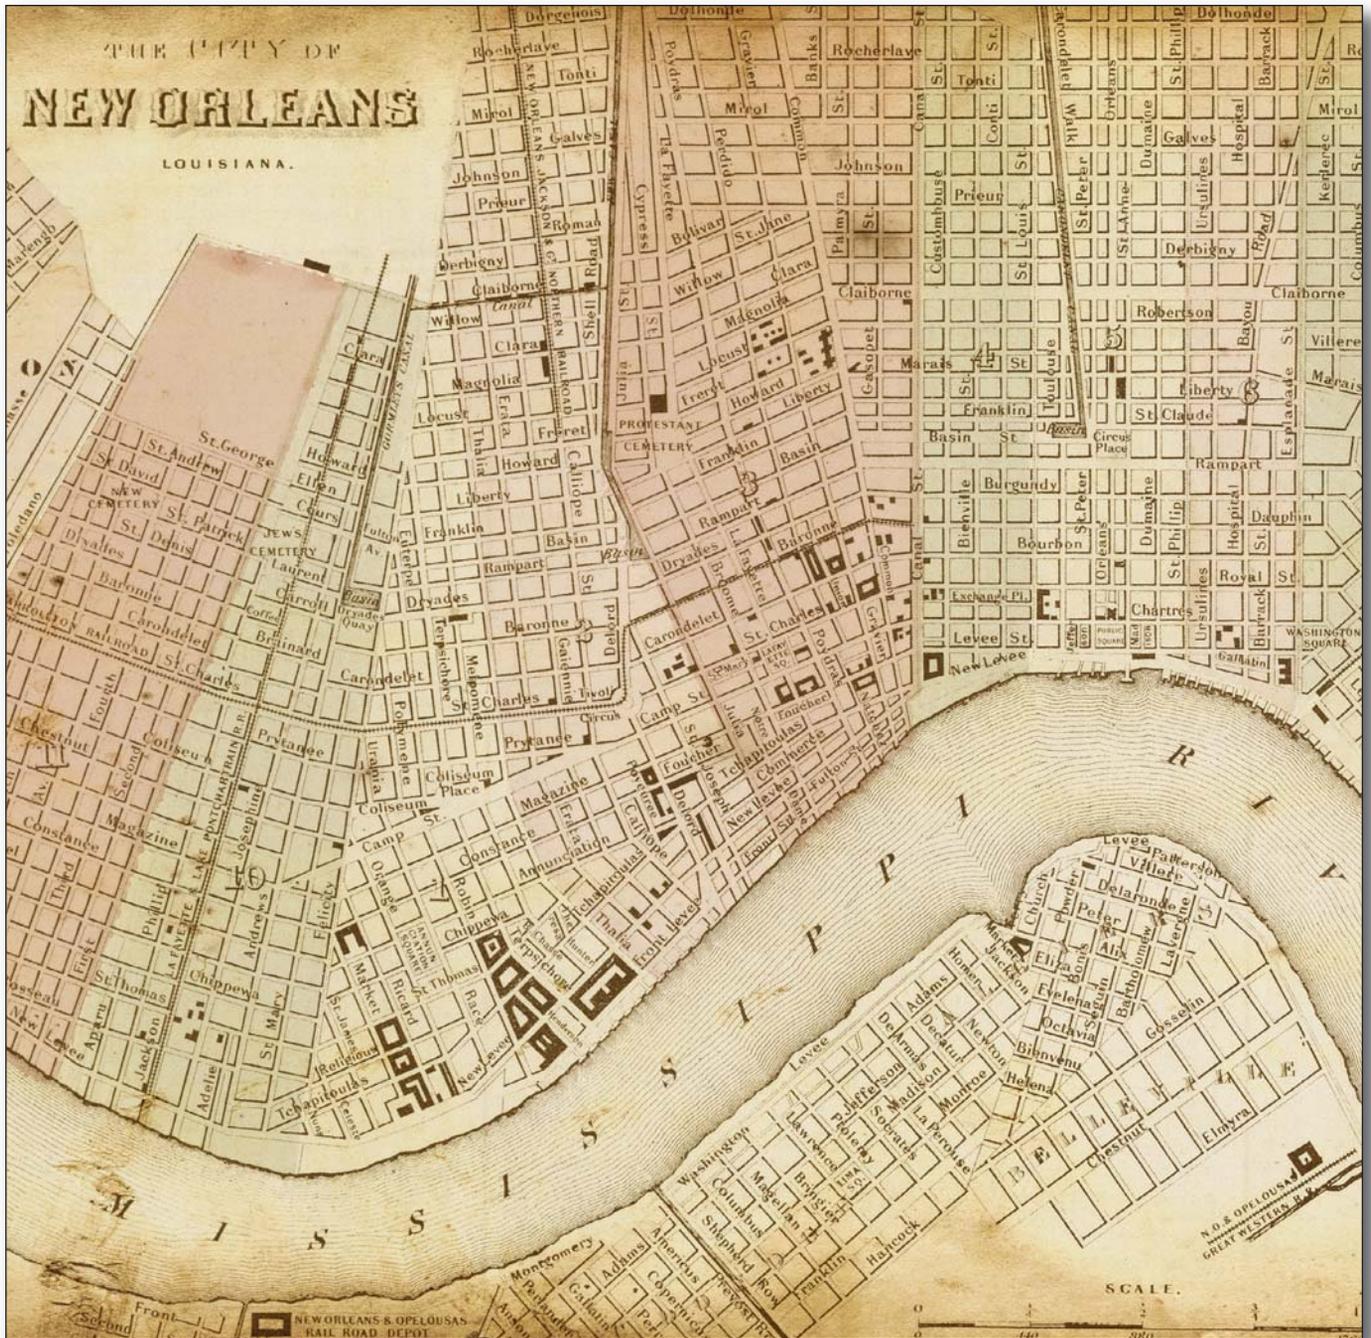
– Sheila McIlraith and Kilian Weinberger

IAAI-18

For 30 years, the Annual Conference on the Innovative Applications of Artificial Intelligence (IAAI-18) has been the premier conference on applied AI research ranging from exciting new potential applications to innovative full deployments of AI technology. IAAI is colocated with AAAI and the paper

presentations and invited talks are coordinated. This coordination enables conference attendees to move seamlessly between conferences that foster interest in applied AI research while keeping track of the latest results of AI research.

As in previous years, IAAI-18 presented papers in three tracks: deployed applications, emerging applications, and challenge problems. The deployed applications track focuses on fielded AI applications that distinguish themselves for their innovative use of AI technology. This track saw significant increase for IAAI-18, with a diverse set of papers: one on improving government human-resource services deployed in three cities; one on improving employee shift schedules with complex constraints used in two types of businesses; one discussing the learnings from deploying sketch technologies in two STEM classes; one on machine-learned revenue forecasting models



AAAI-18 Was Held in the Historic City of New Orleans, Louisiana Located on Banks of the Mississippi River.

used in a major software development company; one on a question-answering system used in an enterprise IT help desk; one on an interactive learning system used to optimize multivariate deployed user interface changes available commercially; and a unique insight into an analytics, analysis, and dynamic adjustment system integrated into the products of several game studios through their major publisher.

The emerging applications track focuses on areas in which AI technology can have a practical impact. For 2018, we had 26 papers in this track — another significant increase for IAAI-18 — covering a broad range of topics such as image analysis in action, providing support to the customer, improving business processes, identifying high value wine, analyzing software and networks, supporting government serv-



AAAI-18 Cochairs Killian Weinberger (left) and Sheila McLraith (right) with Outstanding SPC Member Award Recipient Jörg Hoffmann.

ices, improving utilities and transportation systems, advancing science research, improving communication, advancing entertainment and storytelling, and much more.

The challenge problem track identifies new opportunities and the associated research and development challenges in applying artificial intelligence to real-

world problems. This year we had two papers in this track, one on managing bioengineering complexity and one on the difficulty of batting order setup in cricket.

The conference also featured an outstanding lineup of invited speakers. Zoubin Ghahramani from the University of Cambridge/Uber presented the joint



AAAI's President Rao Kampbampati, President-Elect Yolanda Gil, and Past President Thomas Dietterich at AAAI-18.

AAAI/IAAI-18 talk on probabilistic machine learning and AI and their use in rideshare. Boris Babenko from Orbital Insight presented an IAAI-18 invited talk on leveraging AI and geospatial data to understand the earth at scale. He discussed satellite imagery domain, how it's evolving, and the various advantages and challenges of working with such imagery. He also covered several examples of computer vision modules we have built using deep learning, and some lessons learned. Finally, he spoke about practical considerations of rapidly prototyping and productionizing computer vision and deep learning models.

Finally, in conjunction with *AI Magazine*, Stephen F. Smith (Carnegie Mellon University) received the Robert S. Engelmore Memorial Lecture Award for his

sustained research excellence in constraint-based planning and scheduling technologies, deployment of those technologies to a range of significant real-world problems, and extensive service to the AI community that includes significant outreach to related technical fields.

In his award lecture, Smith discussed the challenging multiagent planning problem implicit in real-time traffic signal control, particularly in urban road networks where (unlike simpler arterial settings) there are competing dominant traffic flows that shift through the day. Further complicating matters, Smith noted, urban environments require attention to multimodal traffic flows (vehicles, pedestrians, bicyclists, buses) that move at different speeds and may be given different priorities. Smith's research group has been evolving an adaptive traffic signal control system to address these challenges, referred to as Surtrac (Scalable Urban Traffic Control). He provided an overview of this overall research effort, noting that Surtrac treats traffic signal control as a decentralized online planning process. In operation, each intersection repeatedly generates and executes (in rolling horizon fashion) signal-timing plans that optimize the movement of currently sensed approaching traffic through the intersection. Each time a new plan is produced (nominally every couple of seconds), the intersection communicates to its downstream neighbors what traffic it expects to send their way (according to the plan), allowing intersections to construct longer horizon plans and achieve coordinated behavior. Smith noted that the initial deployment of Surtrac in the East End of Pittsburgh PA has produced significant performance improvements and that the technology is now being installed in other US cities, and pointed out that recent work in his group has focused on a broader future vision of smart transportation infrastructure where, as vehicles become more connected and more autonomous, the intersection increasingly becomes the gateway to real-time traffic information and navigation intelligence. Current technology development efforts center on use of direct vehicle-to-infrastructure, and pedestrian-to-infrastructure, communication to further enhance mobility, online analysis of traffic flow information for real-time incident detection, and integrated optimization of signal control and route choice decisions.

– G. Michael Youngblood and Karen Myers

EAAI-18

The Symposium on Educational Advances in Artificial Intelligence (EAAI) seeks to advance the AAAI goal of improving the teaching and training of AI practitioners. Organized as an independent symposium within the AAAI conference, EAAI provides the opportunity for researchers, educators, and students to share educational experiences involving AI.



AAAI-18 Outstanding Paper Award Winners Chenjun Xiao and Jincheng Meiers Receive Congratulations from AAAI-18 Cochairs Killian Weinberger and Sheila McIlraith.

EAAI 2018 provided a venue for discussing pedagogical issues, and for sharing resources related to teaching and using AI in education across a variety of curricular levels (K–12 through postgraduate training), with an emphasis on the undergraduate and graduate levels. The symposium showcased ideas for how to teach AI more effectively and highlighted the growing impact of AI in enhancing education. The New and Future AI Educators program, now in its second year, provides fellowships to early career educators and those with an interest in an educational career to enable them to attend EAAI 2018.

EAAI 2018 featured a technical program of long and short talks on refereed papers, including special paper tracks on AI for education, on providing research experiences for undergraduate students, and on best practices for running an AI research group.

To further its aim to foster the sharing of resources and experiences between the attendees, EAAI 2018 also included refereed model AI assignments. Model AI assignments are innovative, engaging and practice-tested course assignments on a variety of AI topics that are ready for use but can also be customized.

EAAI 2018 featured several special events, including an invited talk by Ashok Goel, “Jill Watson, Family, and Friends: Experiments in Building Automated Teaching Assistants.” EAAI also featured Blue Sky Idea talks by the fellowship recipients of the New and Future AI Educators program and two panels. The panel on Next Big Steps in AI for Education looked ahead and identified areas and topics in which AI might play a role in education. The panel on Non-traditional Research Experiences for Undergraduates discussed opportunities for students outside of tradi-



A Lighter Moment During the Presidential Address.

tional NSF-funded REU experiences, such as research-focused courses, senior projects, internships, and competitions like Kaggle.

In his AAAI/EAAI Outstanding Educator Award talk, “Playful AI Education,” Todd Neller took a look more generally at the use of games in AI education. Beginning with a discussion of what makes a game especially good for educational purposes, he then looked at what playful problem domains can bring to educational experiences throughout various AI subfields. The talk concluded with a look forward at goals and challenges for playful AI education.

The symposium closed with presentations and demonstrations by undergraduate students who participated in a workshop to create a demo, game, or other interactive activity that illustrates how artificial intelligence enables autonomous robots.

– Eric Eaton and Michael Wollowski

Eric Eaton is a faculty member in the Computer and Information Science Department at the University of Pennsylvania.

Sheila McIlraith is a professor in the Department of Computer Science at the University of Toronto.

Karen Myers is the program director and principal scientist in SRI International’s Artificial Intelligence Center.

Kilian Weinberger is an associate professor of computer science at Cornell University.

Michael Wollowski is an associate professor of computer science and software engineering at the Rose-Hulman Institute of Technology.

G. Michael Youngblood is an interactive systems researcher, designer, architect, and builder at PARC, a Xerox Company.