Reports of the Workshops of the 31st AAAI Conference on Artificial Intelligence

Monica Anderson, Roman Barták, John S. Brownstein, David L. Buckeridge, Hoda Eldardiry, Christopher Geib, Maria Gini, Aaron Isaksen, Sarah Keren, Robert Laddaga, Viliam Lisy, Rodney Martin, David Martinez, Martin Michalowski, Loizos Michael, Reuth Mirsky, Thanh Nguyen, Michael J. Paul, Enrico Pontelli, Scott Sanner, Arash Shaban-Nejad, Arunesh Sinha, Shirin Sohrabi, Kumar Sricharan, Biplav Srivastava, Mark Stefik, William Streilein, Nathan Sturtevant, Kartik Talamadupula, Michael Thielscher, Julian Togelius, Son Cao Tran, Long Tran-Thanh, Neal Wagner, Byron C. Wallace, Szymon Wilk, Jichen Zhu

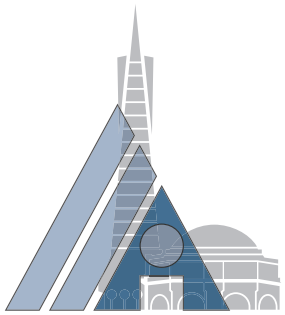

■ *The AAAI-17 workshop program included 17 workshops covering a wide range of topics in AI. Workshops were held Sunday and Monday, February 4–5, 2017, at the Hilton San Francisco Union Square in San Francisco, California, USA.*

The AAAI-17 workshop program included 17 workshops that covered a wide range of topics in AI and were an excellent forum for exploring emerging approaches and task areas, for bridging the gaps between AI and other fields or between subfields of AI, for elucidating the results of exploratory research, or for critiquing existing approaches. Workshops were held Sunday and Monday, February 4–5, 2017, at the Hilton San Francisco Union Square in San Francisco, California, USA.

Workshop participants met and discussed issues with a selected focus — providing an informal setting for active exchange among researchers, developers, and users on topics of current interest. To foster interaction and exchange of ideas, the workshops were kept small, with 25–65 participants. Attendance was sometimes limited to active participants only, but most workshops also allowed general registration by other interested individuals. Most of the workshops were held on a single day.

Of the 17 workshops held, all but 2 (Increasing Diversity in AI and Developing Artificial Intelligence Startup Companies) were included in the AAAI digital library as technical reports.

  ISSN 0738-4602

Organizers of 5 of the AAAI workshops did not submit reports for publication in *AI Magazine*. This report contains summaries of 12 of the workshops that were submitted for publication by organizers. AI Magazine did not receive summaries for the remaining 5 summaries. The summaries included here were edited from the workshop websites or technical reports.

AI and OR for Social Good

The purpose of the AI and OR for Social Good workshop was to explore and promote the application of artificial intelligence (AI) and operations research (OR) for purposes of social good. There has been strong historical interest from both the AI and OR communities on this topic with a burst of AI activity in recent years in topics such as smart grids and optimized transport systems (both as part of a greater computational sustainability effort) while the OR community has long supported areas such as public-sector operations research (PSOR) whose stated objective is doing good with OR.

The workshop placed a special emphasis on bringing together members of the AI and OR communities (notably, the organizing committee consisted of members who overlap with both communities, namely Thomas Dietterich from Oregon State University, Steve Smith from Carnegie Mellon, Pascal Van Hentenryck from the University of Michigan, and Scott Sanner from the University of Toronto) who have been actively involved in addressing challenge problems for social good as well as the AI and OR technologies required to support their solution.

Applications areas targeted for the workshop included but were not limited to sustainable cities, smart government and social services, public service organizations, emergency preparedness, disaster response, public health, and humanitarian programs with problems ranging from data-driven predictive and prescriptive analytics through to logistical optimization. Technical topics targeted included all AI and OR techniques applied to these problems including but not limited to machine learning, constraint optimization and constraint programming, planning and scheduling (under uncertainty), and computational economics.

The workshop program included four invited keynote talks. Milind Tambe from the University of Southern California delivered a talk titled How Can AI be Used for Social Good? Key Techniques, Applications, and Results; Finale Doshi-Velez from Harvard University spoke on AI for Health Care; Daniel Sheldon from the University of Massachusetts Amherst and Mount Holyoke College talked about AI for Ecology and Conservation; and Mark Fox, Center for Social Services Engineering, University of Toronto talked about Measuring Social Good. The workshop also included four sessions of 13 contributed papers covering transport and crime, medicine and health, environment and infrastructure, and social services with contributors representing academic fields such as AI, OR, social work, and health care as well as industry and government participants.

The AI and OR for Social Good workshop was organized by Scott Sanner. This report was written by Scott Sanner. The papers presented at the workshop were published as AAAI Technical Report WS-17-01 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

AI, Ethics, and Society

There is an increasing appetite within and outside AI to hold discussions on AI, ethics, and society. This workshop explored the societal impacts of AI, the impact of AI on jobs and issues like technological unemployment; architectures for ensuring ethical behavior; value alignment in autonomous systems; autonomous agents in the military; autonomous agents in commerce and other domains; measuring progress in AI; and safeguards necessary within AI research. Additional invited sessions discussed some related events (for example, the White House initiative on Preparing for the Future of AI and the New York University meeting on Ethics of AI).

The AI, Ethics, and Society workshop was organized by Toby Walsh. This summary of the workshop was reproduced from the technical report. No report was submitted by the organizer. The papers presented at the workshop were published as AAAI Technical Report WS-17-02 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

AI for Connected and Automated Vehicles

The past decade has witnessed the rapid development of connected and automated vehicles (CAV), which could potentially avoid 90 percent or more of traffic accidents, tremendously mitigate traffic congestion, considerably reduce vehicle energy consumption, and significantly improve the efficiency of roadway usage. However, the existing CAV system is inadequate for the challenges of analyzing large-scale heterogeneous traffic data captured with various vehicle-mounted sensors — cameras, radar, infrared, LIDAR, and making time-critical decisions in complicated driving environments. Solving these two issues goes beyond individual AI techniques, for example, perception, planning, or reasoning, and calls for innovative computing methods that can work in a tightly collaborative manner.

The mission of this workshop was to create a synergy among the AI community — including comput-

er vision, cognition, reasoning, learning, planning, and CAV. The three goals of this event were (1) to identify key AI challenges in CAV systems; (2) to recognize the promising AI solutions to these challenges; and (3) to foster future research in this interdisciplinary subject.

The AI for Connected and Automated Vehicles workshop was organized by Xiaobai (Byran) Liu, Xianfeng (Terry) Yang, Xiaodi Hou, and Mahmoud Tarokh. This summary of the workshop was reproduced from the technical report. No report was submitted by the organizers. The papers presented at the workshop were published as AAAI Technical Report WS-17-03 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15*.

Artificial Intelligence for Cyber Security

The workshop began with a keynote address by Howie Shrobe (MIT CSAIL), titled Leveraging AI for Cyber Security. Shrobe stressed the need for AI in cyber to help address, among other things, the imbalance of effort required by attackers and defenders through the use of symbolic reasoning and statistical machine-learning methods. He advocated an encompassing AI-cyber architecture that predicts then prevents impending attacks, speeds their detection, and contributes to system recovery, thus enabling mission fight-through.

Within the malware and threat session, the first paper presented a game-theoretic approach for deploying configurable honeypots to attract and capture evolving adversary attacks. The method leveraged a multiarmed bandit (MAB) framework to support strategic configuration selection and showed good performance on simulated data. A second paper discussed work in adaptive security, in which a self-aware network leverages a trust model to adapt and avoid cyber-attacks, balancing resource provisioning and attack mitigation.

The challenge problem was presented next. It focused on analytical techniques that leverage cyber-specific domain knowledge to detect cyber threats. A baseline solution was presented and was followed by a discussion on ways to deal with heavily imbalanced cyber data sets.

The machine learning and game theory session began with a paper that presented a knowledge-based expert system for security auditing of cloud platforms that automatically generates human-readable audit reports. The second paper discussed the use of deep, recurrent neural networks to detect anomalous network activity from system logs. The method supports real-time detection and its performance exceeds that of several other well-known techniques on an insider threat data set.

The next paper presented a game-theoretic approach for optimal cyber threat alert prioritization in which an adaptive adversary evades detection by using attacks that cause erroneous alerts. The final paper of the session presented a Stackelberg game model for botnet data exfiltration attacks and corresponding defenses. It proposed greedy heuristics to achieve game equilibrium and generate the best defensive policy.

The afternoon keynote address by Milind Tambe (University of Southern California) was titled Security Games 10 Years after ARMOR: Lessons Learned from Deployed Applications. Tambe illustrated the successful applications of game-theoretic approaches to national security. He emphasized the importance of understanding the user's real-world constraints to facilitate solution adoption.

In the first afternoon session, which featured research leveraging social media data, the initial paper presented research on using machine-learning methods to remove unstructured and irrelevant non-cyber information from social media in order to deliver relevant cyber content to analysts. The next paper discussed new methods that accurately detect spam while simultaneously robustly classifying users that modify their behavior to avoid detection. The final paper of the session presented a framework for automated analysis and categorization of the dark web ecosystem that provides better situational awareness of new content.

The workshop concluded with a panel discussion on the use of AI in real-world operations and included industry panelists Jennelle Bray (LinkedIn), David Burke (Galois), Max Kleiman-Weiner (Diffeo), Sven Krasser (CrowdStrike), and Harold Moss (Akamai). Panelists discussed the need for realistic expectations of the promises of AI for real-world operations — AI tools for security can never be made infallible and the biggest challenge is modeling the rapidly changing environment. The need to share data was discussed with emphasis on the associated legal and practical challenges. Academic-to-industry collaborations were suggested, such as internships and partnerships, as a way to engender trust and permit data sharing with researchers.

This was the second AI for Cyber Security workshop with a 44 percent increase in attendance from the prior year.

The Artificial Intelligence for Cyber Security workshop was organized by William Streilein, Robert Laddaga, David Martinez, Arunesh Sinha, and Neal Wagner. This report was written by the organizers. The papers presented at the workshop were published as AAAI Technical Report WS-17-04 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

AI for Smart Grids and Buildings

The availability of advanced sensing and communication infrastructures, electric monitoring facilities, computational intelligence, widespread use and interest in renewable energy sources, and customer-driven electricity usage, storage, and generation capabilities, have posed the foundations for a robust and dynamic next-generation economic interplay between the demand side: smart buildings, and the supply side: smart power grids.

Three key aspects distinguish this evolving economy from more traditional market forces: (1) information — energy producers and consumers have access to information; (2) exchange — communication is possible on a continuous basis, thus enabling both individual as well as group decision processes; (3) energy can be produced not only by power plants, but also by customers and stored for later use, and (4) customers can employ advanced tactical measures for improving building operations and reducing energy consumption without sacrificing occupant satisfaction.

AI plays a key role in the relationship between the smart grid and smart buildings. New technologies offer infrastructure that provides information to support automated decision making on how to (automatically) adapt production and consumption, optimize costs, waste, and environmental impact, and provide reliability, safety, security, and efficiency. Indeed, several research projects have already developed the view of this ecosystem as a multiagent system, where agents coordinate and negotiate to achieve smart grid and smart building objectives.

The workshop began with an invited talk, by Ole Mengshoel, providing an exciting overview of the applications of machine-learning techniques to sustainability and smart buildings. The workshop provided a predominance of presentations addressing the use of different AI techniques to create the foundations for smart building. He and Fioretto focused on uses of different constraint-optimization methods to support devices that schedule and demand optimization. Jin, Zhan, and Hatalis addressed instead the use of data analytics methods to promote energy conservation, cost reduction, and forecasting of energy production for renewable energy sources. The presentation by Zulas focused on the design of smart homes according to the Solar Decathlon competition. Energy monitoring and energy disaggregation to optimize uses of devices in commercial building was the focus of the presentation by Bansal and Schmidt. Finally, the presentation by Kolev and Johnson analyzed the issues of infrastructure protection.

The workshop was both a model of international and intellectual diversity, with participants from all over the globe, including the United States, the UK, Australia, Germany, Bulgaria, and China. Participants also represented various government agencies and laboratories, industry research laboratories, and academe.

The AI for Smart Grids and Buildings workshop was organized by Rodney Martin, Enrico Pontelli, Son Cao Tran, and Long Tran-Thanh. This report was written by Rodney Martin, Enrico Pontelli, Son Cao Tran, and Long Tran-Thanh. The papers presented at the workshop were published as AAAI Technical Report WS-17-05 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15*.

Computer Poker and Imperfect Information Games

Recent years have brought a substantial progress in research on imperfect information games. There is an active community of researchers focusing on computer poker, which recently computed near optimal strategy for the smallest poker variant commonly played by people and outperformed professional poker players in more complex variants of this game. Game-theoretic models with all sorts of uncertainty and imperfect information have been applied in security domains ranging from protecting critical infrastructure and wildlife to cyber security. Computer agents able to play previously unknown imperfect information games only based on a formal description of its dynamics have been developed.

In this workshop, we aimed to create a forum where researchers studying theoretical and practical aspects of imperfect information games can meet, present their recent results, and discuss their new ideas. Moreover, we tried to facilitate interaction between distinct communities studying various aspects and focusing on various domains in imperfect information games.

Interesting highlights of this year's workshop were invited talks on recent man-machine matches in heads-up no-limit Texas hold'em. Michael Bowling gave a talk about DeepStack, developed by a team at the University of Alberta, which outperformed professional poker players in November and December 2016. Tuomas Sandholm and Noam Brown gave a talk about Libratus, developed at Carnegie Mellon University, which beat heads-up specialists in January 2017. These talks attracted a lot of attention of the participants in other collocated workshops, who joined us for this session.

The main technical program was composed of 11 technical paper presentations and a discussion about the future of the Annual Computer Poker Competition (ACPC). Three of the papers detailed algorithmic advancements and evaluation methods used in the man-machine matches. Other presentations brought novel algorithms for solving imperfect information games; methods for efficient exploitation of static opponents; simplification of complex strategies to make them human learnable; analysis of strategic information revelation in security games; analysis of

the impact of manipulating prices in financial markets; and new theoretical results on games with imperfect recall.

Since achieving statistical significance in the results of ACPC required much more computation than expected, the complete results of the competition were not presented at the workshop this year. The main conclusion of the discussion about the future of ACPC was that there is considerable interest in having a six-player no-limit event. We should try to assess whether it is feasible to achieve statistically significant results with a reasonable amount of computation. However, the event might make sense even as an exhibition without rigorous determination of the winner. Another conclusion from the discussion was that we want to keep the bar for entering the competition reasonably low and motivate new participants to enter the competition.

The Computer Poker and Imperfect Information Games workshop was organized by Viliam Lisy, Michael Thielscher, and Thanh Nguyen. This report was written by Viliam Lisy, Michael Thielscher, and Thanh Nguyen. The papers presented at the workshop were published as AAAI Technical Report WS-17-06 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15*.

Crowdsourcing, Deep Learning, and Artificial Intelligence Agents

Virtual assistants, robots, and other artificial intelligent agents are becoming mainstream in our lives. They manage our calendar, help us navigate to the closest Starbucks, or help us when we're doing online shopping. Thanks to the power of deep learning and cloud computing, machines have been mimicking our brain learning patterns, mainly by ingesting large amounts of data from which they learn how to execute tasks, provide answers, and make decisions. Most of these AI agents are learning to recognize and understand us, even when we talk to them in different dialects and in different languages and are communicating back to us using synthetic voices. Data plays a huge role in the development of these agents and crowdsourcing has been used widely in academe and industry in order not only to help scaling these data needs but also to help developers test the user experience of their virtual assistant apps.

The Crowdsourcing, Deep Learning, and Artificial Intelligence Agents workshop featured five invited talks. Stefano Vegnaduzzo (Integral Ad Science) spoke about challenges and opportunities for crowdsourcing for highly unbalanced classes. Lyle Ungar (University of Pennsylvania) delivered a talk titled Measuring Psychological Traits Using Social Media. Daniel S. Weld (University of Washington) spoke about high-quality crowsources. Daniela Braga (DefinedCrowd), spoke about challenges and opportunities when collecting data for bots. Finally, Ece Kamar's talk concerned troubleshooting AI systems with humans in the loop.

The workshop also included a panel discussion, led by Daniela Braga, with Ece Kamar, Jerome Bellegarda, Lyle Ungar, and Gina-Anne Levow.

The Crowdsourcing, Deep Learning, and Artificial Intelligence Agents workshop was organized by Daniela Braga. This summary of the workshop was reproduced from the technical report and workshop website. No report was submitted by the organizer. The papers presented at the workshop were published as AAAI Technical Report WS-17-07 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15*.

Developing Artificial Intelligence Startup Companies

The objective of the Developing Artificial Intelligence Startup Companies workshop was to bring together members of the AI community with entrepreneurs and those who have been involved in a successful AI startup company, to explore the opportunities and challenges associated with developing successful companies based on artificial intelligence technologies. The workshop considered recent commercial successes in the field and what lessons can be learned.

Talks scheduled for presentation at the workshop included From AI Research to AI Startup by Markus Fromherz (SK Telecom Americas); Lessons from Production AI by Fritz Henkel (ASAPP); Accelerating AI: Lessons from Backing 40 AI Companies a Year in Hardware, Life Sciences, Software, and Food by Sean O'Sullivan (SOSV); The Internet of Trees by Enda Keane (TreeMetrics); and AI for Enterprise Procurement by Alan Holland (Keelvar).

The Developing Artificial Intelligence Startup Companies workshop was organized by Barry O'Sullivan, Markus Fromherz, and Wayne Murphy. This summary of the workshop was reproduced from the workshop website. No report was submitted by the organizers and no technical report was issued.

Distributed Machine Learning

With the fast development of machine learning (especially deep learning) and cloud computing, it has become a trend to train machine-learning models in a distributed manner on a cluster of machines. In recent years, much exciting progress has been made along this direction, with quite a few papers published, and several open-source projects populated. For example, distributed machine-learning tools such as Petuum, TensorFlow, and DMTK have been developed; parallel learning algorithms such as LightLDA, parallel logistic regression, XGBoost, and PV-Tree have been proposed; and convergence theories for

both synchronous and asynchronous parallelization have been established. However, there are also many open issues in this field, such as how to select an appropriate infrastructure and parallelization mechanism given the application and system configuration, why many papers report linear speedups, but when the accuracy on real-world workloads, the practical speed-up is far smaller, why parallelization mechanisms with similar convergence rates perform so differently in practice, and how one conducts proper comparison and evaluation for distributed machine learning (for example, benchmark, criteria, system configurations, and baselines).

Without answers to these important questions, people can hardly be confident in wide adoption of distributed machine learning in real applications. The hope for this workshop was to provide the community with deep insights and to substantially push the frontier of distributed machine learning.

Invited talks were delivered by Alex Smola (Amazon), Joseph E. Gonzales (University of California, Berkeley), Xiangrui Meng (Databricks), and Christopher Ré (Stanford University). The Distributed Machine Learning workshop was organized by Tie-Yan Liu, James Kwok, and Chih-Jen Lin. This summary of the workshop was reproduced from the technical report. No report was submitted by the organizers. The papers presented at the workshop were published as AAAI Technical Report WS-17-08 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15*.

Health Intelligence

Population health intelligence includes a set of activities to extract, capture, and analyze multidimensional socioeconomic, behavioral, environmental, and health data to support decision making to improve the health of different populations. Advances in artificial intelligence tools and techniques and Internet technologies are dramatically changing the ways that scientists collect data and how people interact with each other and with their environment. Moreover, the Internet is increasingly used to collect, analyze, and monitor health-related reports and activities and to facilitate health-promotion programs and preventive interventions. In additon, to tackle and overcome several issues in personalized health care, information technology will need to evolve to improve communication, collaboration, and teamwork among patients, their families, healthcare communities, and care teams involving practitioners from different fields and specialties.

This Health Intelligence workshop follows the success of the earlier workshops held in conjunction with the 27th, 28th, 29th, and 30th AAAI Conferences on Artificial Intelligence. This joint workshop brought together a wide range of particpants (about 50 registrants) from the multidiciplinary field of medical and health informatics. Participants were interested in the theory and practice of computational models of web-based public health intelligence as well as personalized health-care delivery. The papers and demonstrations presented at the workshop covered a broad range of disciplines within artificial intelligence including knowledge representation, machine learning, natural language processing, pattern recognition, digital imaging, and online social media analytics. From an application perspective, presentations addressed topics in epidemiology, environmental and public health informatics, disease surveillance and diagnosis, patient participation, health behavior monitoring, and disaster management.

The workshop also included four invited talks. Rumi Chunara (Global Institute of Public Health, New York University) gave a presentation on the use of unstructured data in population health. Urmimala Sarkar, MD (University of California San Francisco and San Francisco General Hospital) described values of social media in health IRL applications. Mor Peleg (University of Haifa) presented her findings from the MobiGuide Project on how to promote patients' engangement in their health-care decision-making process. John H. Holmes (University of Pennsylvania Hospital) also gave an insightful presentation on AI-driven approaches to data integration.

To promote open debate and exchange of opinion among participants, the workshop held a panel discussion moderated by David L. Buckeridge and included Deborah L. McGuinness (Rensselaer Polytechnic Institute), Rumi Chunara (New York University), and José Luis Ambite (University of Southern California). The major theme of the panel was to discuss the synergy between precision health for individuals and populations.

The Health Intelligence joint workshop was organized by Arash Shaban-Nejad and Martin Michalowski. This report was written by Arash Shaban-Nejad, Martin Michalowski, David L. Buckeridge, Byron C. Wallace, Michael J. Paul, Szymon Wilk, and John S. Brownstein. The papers presented at the workshop were published as AAAI Technical Report WS-17-09 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Human-Aware Artificial Intelligence

As AI techniques and systems come into increasing contact with humans, and into the public consciousness at large, various research issues surrounding such interactions have come to the fore. Specifically, a key movement that is underway in the AI community and the world of technology at large concerns the notion of humans and machines (AI systems) teaming up together to understand data and take deci-

sions. The key premise of this workshop was based on the idea that augmented intelligence — that is, teams and systems that combine the skills of humans and AI techniques — can achieve better performance than either alone. However, to create such systems with augmented intelligence, humans must be accommodated as first-class citizens in the decision-making loop of existing AI systems. Far too often, traditional AI systems have tended to exclude humans (and the problems that accompany interaction with them) and have instead focused on producing optimal artifacts that stand no significant chance of working in the real world.

To address this issue and produce truly human-aware artificial intelligence, systems must try to solve the interaction issues that accompany each unique application domain. These interaction issues may broadly be divided into extraction (or interpretation) challenges and presentation (or steering) challenges. Extraction challenges deal with understanding human input, whether that be in the form of knowledge, or in the form of specific directives and goals to achieve. Presentation challenges deal with questions of how to present the system's outputs to the team and solicit feedback. However, the specific interaction issues may differ significantly depending on the application being addressed.

The workshop brought together researchers from various subfields of AI, including automated planning, multiagent systems, game theory, execution monitoring, game playing, decision theory, machine learning, and integrated decision architectures. The participants also represented a wide swathe of AI stakeholders from academe, industrial research labs, technology startups, policy bodies, and the government.

The full-day workshop was split into two halves, and two main themes dominated the discourse through the day. The first was a retrospective of past human-aware AI systems and their contributions and shortcomings. The invited talks that touched on this theme were tied together by a panel on the various challenges in involving humans in AI systems — one major issue that was identified concerned the various models of a single domain (the machine's and the human's, but also each participant's model of the other), and the impedance mismatch between these various models. Other issues that came up were explaining the machine's decision process to the human, and the need to have transparent and reasonably stable and consistent behavior from the AI system.

The second theme revolved around the future of such human-machine augmented domains and systems and various ethical and regulatory questions surrounding this. Issues that were raised included the need for human-aware AI when considering social and societal problems (AI for social good), ethical concerns and professional codes for deployed and consumer-facing AI, and the building of trust between AI and the humans who either consume its decisions or are affected by them.

The workshop participants expressed their appreciation at the existence of a large and diverse community that was interested in these issues and wished to continue discussion and collaboration under the broad umbrella of human-aware artificial intelligence. Specific outcomes from the day-long workshop include the creation of a common forum for the exchange of news and ideas on human-aware artificial intelligence and a proposal to run a dedicated special track at the upcoming premier AI conferences on the issue.

The Human-Aware Artificial Intelligence workshop was organized by Kartik Talamadupula, Shirin Sohrabi, Biplav Srivastava, and Loizos Michael. This report was written by Kartik Talamadupula, Shirin Sohrabi, Biplav Srivastava, and Loizos Michael. The papers presented at the workshop were published as AAAI Technical Report WS-17-10 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Human-Machine Collaborative Learning

Early work in AI relied on extensive knowledge engineering from human domain experts. This has given way to modern machine-learning-based AI, which relies on vast numbers of training examples to learn from instead. These AI systems also tend typically to operate in a stand-alone fashion. In this context, a new field of collaborative learning between humans and machines is emerging, where the goal is to enable AI to learn with assistance from domain experts in place of requiring extensive knowledge engineering or training examples, and participating in teams in place of stand-alone operation, which allows the AI to focus on more systematic tasks while allowing humans to address more open-ended problems. The goal of this symposium was to investigate the current state of collaborative learning and identify research problems that will need to be addressed to enable efficient collaborative learning for real-world problems.

The workshop brought together researchers from a variety of fields including machine learning, statistics, deep learning, reinforcement learning, human-computer interfaces, decision-support, and planning. The first major theme of papers presented at the symposium was in the space of advanced methods to seek and incorporate knowledge from domain experts. Techniques in this space included novel forms of apprenticeship learning, incorporating knowledge into relational models, learning options automatically, and understanding incomplete instructions. The other theme was in the space of AI designed to support humans in reducing the complexity of a task,

primarily through the use of dimension reduction and visualization based approaches.

The workshop helped highlight the progress that was being made in both individual directions: (1) efficiently incorporating knowledge from domain experts to solve a task, and (2) assisting domain experts with information to help the domain experts solve a task. Perhaps more importantly, the workshop also helped the participants recognize that there was limited work and significant, but as of yet unexplored, potential in the space of truly collaborative learning, which would involve directions (1) and (2) both occurring jointly, with different subtasks being split between the AI agent and the domain expert. We are hopeful that this workshop on collaborative learning will help kick-start this field, and lead to significant breakthroughs in the coming years in critical areas for enabling collaborative learning including explainable AI and human-machine communication, understanding of optimal decomposition of tasks between humans and machines, and management of swarms of humans/machines.

The Human-Machine Collaborative Learning workshop was organized by Hoda Eldardiry. It was led by PARC researchers Hoda Eldardiry, Kumar Srcicharan, and Mark Stefik. This report was written by Hoda Eldardiry, Kumar Srcicharan, and Mark Stefik. The papers presented at the workshop were published as AAAI Technical Report WS-17-11 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Increasing Diversity in Artificial Intelligence

Participation in AI by groups traditionally underrepresented in computer science is a fraction of what is needed to have a workforce that reflects the diversity in the society. The decrease in the number of women and members of other underrepresented groups in AI is especially worrisome, because AI over the years had enjoyed a larger representation of women compared to other areas of computer science.

The goal of the Increasing Diversity in Artificial Intelligence workshop was to discuss the importance of diversity in AI in multiple senses. The workshop started as an opportunity to get together as a community to talk about underrepresentation of many groups in AI and how to address it, but it was broadened to discuss diversity in a broader sense, namely also diversity of areas of research within AI, diversity of the methods used in AI, and diversity of communities within AI.

At the workshop, invited speakers presented brief overviews of different areas of research within AI, ranging from hybrid systems that combine the strengths of machine and human intelligence to game theory and its use in conjunction with AI-inspired frameworks and models, the future of automated programming, and the challenges in the study of large-scale social networks. The talks provided a gentle introduction to many exciting research areas, methods, and open research problems.

The workshop included a talk on what's hot in AI ethics, a topic of growing importance, and a career panel with a mix of academic and industry speakers. As part of the discussion with panelists, more diversity issues that are often neglected were brought up, such as age diversity and socioeconomic diversity. Many students from underrepresented groups graduate heavily in debt. They cannot afford to take time off after school to travel. They might never have had a mentor who encouraged them to think ambitiously. In some cultures, ambition is not socially acceptable. What happens when people change careers, going back to college for a new undergraduate degree or for a graduate degree after years in the work force? Can they expect to have the same career opportunities as traditional students or are some career paths blocked off for them? What happens when students with a different socioeconomic background get into PhD programs? Will the lack of financial backing make it difficult or impossible for them to pursue research and networking opportunities that require financial exposure?

The program of the day was concluded by Marie desJardins, who talked about how to balance the multiple tasks that academe expects from assistant professors and how to enjoy an academic or research career in AI.

Funding from CRA-W, under the organization's Discipline Specific Workshop program, was awarded to student and postdoc participants, enabling them to attend the workshop and the conference. Adele Howe (Colorado State University), who was a cochair of the workshop, passed away shortly before the workshop. Adele was a tireless leader in promoting access and diversity in the AI community. She is greatly missed.

The Increasing Diversity in Artificial Intelligence workshop was organized by Monica Anderson, Amy Greenwald, Judy Goldsmith, and Adele Howe. This report was written by Maria Gini and Monica Anderson. No technical report was issued.

Knowledge-Based Techniques for Problem Solving and Reasoning

Despite recent attempts in various subareas of AI to integrate technologies to solve complex problems such as autonomous cars, there are still gaps between research communities that prevent efficient transfer of knowledge. For example, knowledge representation techniques focus on formal semantics and flexibility of modeling frameworks and put less emphasis on actual problem solving that requires efficient tools. Other communities such as planning and

search put emphasis on efficiency of problem solving, but give less attention to how the real problem is modeled, the connection between modeling and efficiency of problem solving, and the capability of the models to support other important features like plan revision and adaptation. The Knowledge-Based Techniques for Problem Solving and Reasoning workshop attempted to bridge these particular communities with the goal of exchanging information leading to more efficient problem solving starting with the problem requirements and finishing with the solved problem.

The workshop was driven by the idea that a good model is necessary for efficient problem solving and hence formal problem modeling should be studied more tightly with problem solving. In some sense, the workshop complemented the current trend of using model-free approaches such as deep learning to solve every type of AI problem. The workshop attempted to answer questions on relations between the formal models and problem solving. How do the formal models relate to efficiency of problem solving? How do various modeling frameworks compare from the perspective of problem solving? How can the model be acquired? How can the model be verified and validated? How can the formal model be reformulated to get an efficiently solvable model? How can the solution be checked with respect to the model? How does the model evolve in time? How can the model support solution revisions at execution time?

This was the second workshop in a short series that started at the 2016 International Joint Conference on Artificial Intelligence. It follows the tradition of more specific workshops such as ICAPS workshops on knowledge engineering for planning and scheduling. The workshop attracted researchers from several areas of AI as the topics of the workshop are naturally broad. This was also reflected in presented papers. Four major themes can be identified in papers presented during the workshop. The first theme was planning, where the presented papers discussed, for example, how to learn axioms to improve efficiency of planning and how to deduce new information about only partially specified initial states. The second theme was learning causal knowledge from a sequence of events, for example, learning that thunder follows lightning. The third theme dealt with knowledge graphs, namely how to create them from unstructured text and how to transfer knowledge between knowledge graphs. The last theme was related to various applications such as detecting the number of people using simple sensors in smart homes and knowledge-based classification of galaxies.

The common theme of presented papers was exploiting knowledge for more efficient solving of various problems. The majority of presentations dealt with how to obtain such knowledge; for example, missing information can be extracted from a general knowledge about the domain or transferred from a different model. Some presentations showed how to exploit knowledge in solving specific problems. The major outcome of the workshop was that there is definitely a need for AI research subcommunities to meet and to discuss the topics on the border between the subareas.

The Knowledge-Based Techniques for Problem Solving and Reasoning workshop was organized by Roman Barták, Thomas Leo McCluskey, and Enrico Pontelli. This report was written by Roman Barták. The papers presented at the workshop were published as AAAI Technical Report WS-17-12 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Plan, Activity, and Intent Recognition

Work in the plan, activity, and intent recognition community uses many different representations, which were reflected in the papers presented to the Plan, Activity, and Intent Recognition workshop — activity recognition uses raw data as input (sometimes accompanied with a plan library or planning domain), and labels an output sequence of observations (Deep LSTM-Based Goal Recognition Models for Open-World Digital Games by Wookhee Min, Bradford Mott, Jonathan Row, and James Lester). Goal and intent recognition take in a partial or full sequence of observations and a planning domain as input and outputs the intended goal of the agent (Plan Optimality Monitoring Using Landmarks and Planning Heuristics by Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi). Plan recognition uses a partial sequence of observations and a plan library or planning domain as input and outputs either a sequence of future steps or a hierarchical plan (Partial Observability in Grammar-Based Plan Recognition by Christopher Geib and Robert Goldman). The recent advent of work on plan recognition as planning takes as input a planning domain and a set of possible goals and selects one of the goals (An AI Planning-Based Approach to the Multiagent Plan-Recognition Problem by Maayan Shvo, Shirin Sohrabi, and Sheila A. McIlraith). Although all of these problems have much in common, there is no single standard representation to allow comparison of the work of the kind used by the 1998 International Planning Competition. This year, in addition to presenting new work, the Plan, Activity, and Intent Recognition workshop deliberately emphasized the discussion of establishing test suites, benchmarks, and challenge problems to address this need.

A comparison of works in the plan-recognition community and selected domains makes it evident that there are almost no common data sets or even metrics for evaluation of different recognizers. The situation becomes even worse when trying to compare these works to activity and intent recognition.

In the workshop, we discussed establishing common grounds for future comparisons and evaluations and the need to create standard representations and standard problems. The workshop participants took part in a discussion about an initial proposal for a standard domain representation and agreed that in the coming year we would develop (1) a forum for the standardization project in which participants will be able to influence deci-

sions regarding representations, task scopes, and other factors; (2) a standard representation that can encapsulate both plan libraries and planning domains; and (3) a list of four to five plan, activity, and intent recognition tasks with standard domains that will be presented next year for comparison.

Plan, activity, and intent recognition plays a crucial role in a wide variety of applications including personal intelligent assistants, assistive technology in health and smart environments, intelligent human-computer interfaces, natural language and speech dialogue management, computer and network security, coordination in robots and software agents. Researchers who investigate these various tasks come to realize the importance of inferencing the agent's behaviors and the need for explainable AI.

These are exciting times in our research community and the Plan, Activity, and Intent Recognition workshop provided an excellent opportunity to share results, discuss potential for integration and collaboration, and accelerate efforts at standardization

The Plan, Activity, and Intent Recognition workshop was organized by Reuth Mirsky, Sarah Keren, and Christopher Geib. This report was written by Reuth Mirsky, Christopher Geib, and Sarah Keren. The papers presented at the workshop were published as AAAI Technical Report WS-17-13 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Symbolic Inference and Optimization

Symbolic methods enjoy a long and distinguished history in AI. While the last two decades have seen major advances in probabilistic modeling, data management, data fusion, and data-driven learning, much of this work assumes fairly low-level representations, tailored toward a specific application. It is now recognized that formal languages, and their symbolic underpinnings, can enable descriptive clarity, reusability, and interpretability, thereby furthering the applicability and impact of AI technology. Motivated by recent breakthroughs in formal representations and symbolic techniques for inference and optimization, the workshop brought together researchers working in the areas of probabilistic modeling, weighted model counting, symbolic logic, planning, and numerical optimization in a push to identify common theoretical and algorithmic ideas across these fields.

The discussions of the workshop were centered on two major themes. The first theme was probabilistic modeling and inference with circuit representations. In invited talks, Stefano Ermon (Stanford University) gave an overview of the theory of Fourier-space representations for probability distributions, which have recently been applied with great success to probabilistic tasks. Guy Van den Broeck (University of California, Los Angeles) discussed another highly successful circuit representation that admits tractable inference, namely probabilistic sentential decision diagrams. In essence, these two representations present different ways of exploiting the idea of recursive conditioning to achieve tractability. Another take on the principle manifests itself in the concept of AND/OR search spaces, presented by Rina Dechter in her talk on combinatorial search for probabilistic inference tasks. This topic was also reflected in the papers of the workshop, which presented contributions in terms of inference algorithms and richer representations.

The second major discussion of the workshop revolved around extending symbolic representations beyond standard inference tasks or doing inference by nonstandard means. Of special interest was the idea of solving problems traditionally considered within the realm of numerical approximation by means of symbolic inference. This idea was illustrated by Steven Diamond (Stanford) in a keynote discussing convex optimization with abstract (symbolic) linear operators. Next to convex optimization, contributed papers addressed the problems of symbolic stochastic planning and shortest path problems. The converse idea, symbolic inference by means of numerical optimization, was also discussed during the workshop.

The workshop participants discussed how even though symbolic approaches for inference are as old as AI itself, there are many unexplored opportunities to apply them, especially within machine learning and numerical approximation tasks. It was remarked how the workshop was useful in distilling the core ideas used throughout the field and extending them to new horizons, and the participants expressed interest in attending future workshops on the topic.

The Symbolic Inference and Optimization workshop was organized by Scott Sanner. This report was written by Scott Sanner. The papers presented at the workshop were published as AAAI Technical Report WS-17-14 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

What's Next for AI in Games?

As we look to future work on AI in games, this workshop encouraged communication and sharing of ideas between people currently working on AI research specifically focused on games, working on core AI research that could be applied to games, and working in the games industry eager to apply advances in AI for the design, analysis, and play of games. Although much public focus on game AI is on improving the strength of opponents and the realism of nonplayer characters, there were many other uses for AI in games and game development discussed in the workshop. The workshop had 11 full length paper presentations, 6 future challenge presentations, and a panel on AI in the games industry.

Many of the topics presented elsewhere in the wider AAAI conference were presented at the workshop, focusing on their applications to games. Deep learning and machine learning were hot topics, and the workshop included papers from across the globe on deep reinforcement learning agents for playing Doom, efficient deep reinforcement learning using multifocus attention networks, using anthropo-

morphism game theory to set weights for deep neural networks, how to detect affordances from game play traces, and an analysis of action-space reward structures in games. Natural language processing was also discussed, in the context of automatically extracting story graphs from natural language stories as well as a short paper discussing the gap between computational narrative and natural language processing.

A deep study of games and game strategies, enabled and inspired by AI techniques, was another popular topic. Presentations included a formal analysis of depth in strategic games, using AI agents to evaluate modern board games using search-driven play testing and methods for discovering efficient strategies for Minesweeper. Procedural content generation, a popular area of games research where game content such as levels, maps, and characters are created algorithmically, was covered in the contexts of large-scale search, cognitively grounded procedural content generation, and generating content for game spectators as opposed to game players.

The impact of AI and games on players, modern humans, games research, and humanity's history was discussed in several presentations. These varied topics included ethical considerations for player modeling, an anthropological study of ancient games as a new AI frontier, the ability to better match players for more positive game experiences, and a discussion of how reliance on certain types of AI methods may be limiting our research.

We concluded the event with a panel on What's Next for AI in the Games Industry, with highly experienced panelists Robin Hunicke (Funomena, ThatGameCompany, University of California, Santa Cruz), Frank Lantz (New York University Game Center, Area/Code Entertainment, Zynga New York), Ben Weber (Twitch, EA), and Alexander Zook (Riot Games, Blizzard Entertainment), moderated by Aaron Isaksen (Indie Fund, Fig, New York University Game Innovation Lab). The panel addressed many interesting topics including how artificial intelligence will change the tools and methods by which games are made, new types of games and game genres enabled by AI, and the impact of AI on the game industry and game education.

The presentations and panel were live streamed, and a recording is available online at www.youtube.com/watch?v=7b2jCiyR1cg.

The What's Next for AI in Games? workshop was organized by Nathan R. Sturtevant, Aaron Isaksen, Julian Togelius, and Jichen Zhu. This report was written by Aaron Isaksen, Nathan Sturtevant, Julian Togelius, and Jichen Zhu. The papers presented at the workshop were published as AAAI Technical Report WS-17-15 in the AAAI Digital Library and included in *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports WS-17-01 – WS-17-15* compilation.

Monica Anderson is an associate professor at the University of Alabama.

Roman Barták is a professor at Charles University, Czech Republic.

John S. Brownstein is affiliated with Boston Children's Hospital, Harvard University.

David L. Buckeridge is affiliated with McGill University.

Hoda Eldardiry leads the machine-learning group at Palo Alto Research Center.

Christopher Geib is an associate professor at Drexel University.

Maria Gini is a professor at the University of Minnesota.

Aaron Isaksen is a PhD candidate at New York University.

Sarah Keren is affiliated with Technion University.

Robert Laddaga is a research professor at Vanderbilt University.

Viliam Lisy is an assistant professor at the Czech Technical University in Prague.

Rodney Martin is a research scientist at NASA Ames Research Center.

David R. Martinez is an associate division head at the MIT Lincoln Laboratory.

Martin Michalowski is affiliated with the University of Ottawa.

Loizos Michael is an assistant professor at the Open University of Cyprus.

Reuth Mirsky is a PhD candidate at Ben-Gurion University.

Thanh Nguyen is a postdoctoral fellow at University of Michigan.

Michael J. Paul is affiliated with the University of Colorado Boulder.

Enrico Pontelli is a professor at New Mexico State University.

Scott Sanner is an assistant professor at the University of Toronto, Canada.

Arash Shaban-Nejad is affiliated with the University of Tennessee Health Science Center, Oak Ridge National Laboratory.

Arunesh Sinha is an assistant research scientist at the University of Michigan.

Shirin Sohrabi is a research staff member at IBM T. J. Watson Research Center.

Kumar Sricharan is a senior research scientist at Palo Alto Research Center.

Biplav Srivastava is a research staff member at IBM T. J. Watson Research Center.

Mark Stefik leads the human-machine collaboration group at Palo Alto Research Center.

William W. Streilein is a group leader at the MIT Lincoln Laboratory.

Nathan Sturtevant is an associate professor at the University of Denver.

Kartik Talamadupula is a research staff member at IBM T. J. Watson Research Center.

Michael Thielscher is a professor at the University of New South Wales.

Julian Togelius is an associate professor at New York University.

Son Cao Tran is a professor at New Mexico State University.

Long Tran-Thanh is a lecturer at the University of Southampton.

Neal Wagner is a technical staff member at the MIT Lincoln Laboratory.

Byron C. Wallace is affiliated with Northeastern University.

Szymon Wilk is affiliated with Poznan University of Technology, Poznan, Poland.

Jichen Zhu is an associate professor at Drexel University.