

Reports of the 2016 AAAI Workshop Program

Stefano V. Albrecht, Bruno Bouchard, John S. Brownstein, David L. Buckeridge, Cornelia Caragea, Kevin M. Carter, Adnan Darwiche, Blaz Fortuna, Yannick Francillette, Sébastien Gaboury, C. Lee Giles, Marko Grobelnik, Estevam Hruschka, Jeffrey O. Kephart, Parisa Kordjamshidi, Viliam Lisy, Daniele Magazzeni, Joao Marques-Silva, Pierre Marquis, David Martinez, Martin Michalowski, Zeinab Noorian, Enrico Pontelli, Alex Rogers, Stephanie Rosenthal, Dan Roth, Scott Sanner, Arash Shaban-Nejad, Arunesh Sinha, Tran Cao Son, William Streilein, Sylvie Thiebaux, Byron C. Wallace, Toby Walsh, Michael Witbroc, Jie Zhang

■ *The Workshop Program of the Association for the Advancement of Artificial Intelligence's Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16) was held at the beginning of the conference, February 12–13, 2016. Workshop participants met and discussed issues with a selected focus, and the workshop provided an informal setting for active exchange among researchers, developers, and users on topics of current interest. The AAAI-16 workshops 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. The 15 workshops held at AAAI-16 were Artificial Intelligence Applied to Assistive Technologies and Smart Environments (WS-16-01), AI, Ethics, and Society (WS-16-02), Artificial Intelligence for Cyber Security (WS-16-03), Artificial Intelligence for Smart Grids and Smart Buildings (WS-16-04), Beyond NP (WS-16-05), Computer Poker and Imperfect Information Games (WS-16-06), Declarative Learning Based Programming (WS-16-07), Expanding the Boundaries of Health Informatics Using AI (WS-16-08), Incentives and Trust in Electronic Communities (WS-16-09), Knowledge Extraction from Text (WS-16-10), Multiagent Interaction Without Prior Coordination (WS-16-11), Planning for Hybrid Systems (WS-16-12), Scholarly Big Data: AI Perspectives, Challenges, and Ideas (WS-16-13), Symbiotic Cognitive Systems (WS-16-14), and World Wide Web and Population Health Intelligence (WS-16-15).*

Artificial Intelligence Applied to Assistive Technologies and Smart Environments

Smart environments have the potential to enhance the quality of life of people by giving assistance in the activities of daily life. These systems are particularly interesting for impaired and frail people because they can improve their autonomy and reduce the need for caregivers. The development of technologies to improve the building of assistive systems has increased this last decade. However, despite this growing interest toward these technologies, no real adoption tendency has been observed yet. Indeed, impairments and particularities of users are so diverse that implementing solutions that are mandatory for users' well-being represent one of the major challenges in terms of universal design. The goal of this workshop was to investigate new solutions to scientific problems occurring in various topics related to artificial intelligence applied in the domain of impaired people assistance.

The 2016 AAAI workshop brought together academic and industrial researchers from several subfields of AI. One main theme of papers presented at the workshop was activity recognition. In fact, this theme is one major challenge in providing good assistance. Indeed, assistive systems need to get accurate information about monitored people. This can be considered as the first step. Several papers provided models and analysis about performing an efficiency detection from different kinds of data.

Another major theme was actions performed by assistive devices. We can divide the papers on this theme into three categories. Some papers presented methods to improve the efficiency of the decision process. Others focused on adaptation of actions to address the needs of users. Finally, some papers studied the ethical issues in autonomous systems.

The workshop participants discussed how assistive technologies can benefit from innovation in domains such as the internet of things. However, participants shared the objective of developing methods to reduce the costs and increase the efficiency of assistive technology and agreed that they would like to attend a future edition of the workshop.

Bruno Bouchard and Sébastien Gaboury served as cochairs of this workshop. The papers of the workshop were published as technical report WS-16-01 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence*.

AI, Ethics, and Society

This workshop focused on the ethical and societal implications of building AI systems. It is a response to the increasing appetite from both within and outside the AI research community for such discussions.

The topics addressed within the workshop included the future of AI; AI as a threat to or savior for humanity; mechanisms to ensure moral behaviors in AI systems; safeguards necessary within AI research; and the impact of AI on work and other aspects of our lives. The workshop was divided into two parts. In the morning, there were seven contributed talks and six posters. The titles of the talks give a good idea of the range of topics covered: Human-like Morality and Ethics for Robots; Patience Is Not a Virtue: AI and the Design of Ethical Systems; Quantilizers: A Safer Alternative to Maximizers for Limited Optimization; Why the Technological Singularity May Never Happen; Modeling Progress in AI; Taxonomy of Pathways to Dangerous Artificial Intelligence; and Reinforcement Learning as a Framework for Ethical Decision Making.

In the afternoon, there were a dozen

or so short talks from recipients of the Future of Life grants on AI safety. These grants are funded by a \$10 million donation from Elon Musk to promote safe and beneficial AI research. The workshop finished with a panel discussion on the most promising research directions for keeping AI beneficial. The panelists were Stuart Russell, Vincent Conitzer, David Parkes, Percy Liang, Stefano Ermon, and Benjamin Rubinstein.

Special thanks go to Richard Mallah who, in his capacity as part of the organizing committee, put together the program for the afternoon. Toby Walsh served as chair of the workshop. The papers of the workshop were published as technical report WS-16-02 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Artificial Intelligence for Cyber Security

This workshop focused on research and applications of artificial intelligence to cyber security, including machine learning, game theory, natural language processing, knowledge representation, and automated and assistive reasoning. Talks throughout the day emphasized cyber systems and research on techniques to enable resilience in cyber security systems augmented by human-machine interactions. The workshop began with a keynote speech by George Cybenko (Dartmouth College) on Cyber Security Challenges Amenable to AI. Cybenko suggested a need to address the AI challenges across the full spectrum, including deterrence, protection, intrusion detection, adaptation, and recovery. AI can help, since the growth in cyber vulnerabilities is, in many cases, exploiting user behavior (for example, phishing as an entry point). The role of AI is in predicting patterns of behavior to address potential vulnerabilities.

On the topic of malware, papers presented included techniques employing semantic meaning to identify clusters of malware families, showing good true positives versus low false positives in identifying clusters by using integer linear programming. Additionally, an approach was presented using Bayesian

networks to determine malware lineage present in a directed acyclic graph.

Two papers followed focusing on human-machine interaction (HMI). One paper integrated prior knowledge of actor behavior and leveraged machine-learning techniques to work on a reduced space. This enabled the application of rules to represent the actor behavior and better accuracy in final cyber attribution. The second paper leveraged the Levenshtein distance to measure cognitive burden during password creation, usage, and recollection, in order to develop strong password recommendation tools.

There were three position papers within the workshop that were shorter presentations on a topic of interest, including the need to augment cyber security tools by incorporating learning from human behavior; recommendations for a multiarmed bandit approach to quickly structure and configure honeypots; and using game theory for incentivizing participants in sharing information while incorporating timely updates to improve the incentives.

The afternoon keynote from Robert Laddaga (Vanderbilt University) was titled "AI and the Future of Cyber Security." Laddaga argued that the use of AI technologies in cyber security will inevitably improve both defenders and attackers. He stressed the need for active methods to detect adversary activity, protect cyber systems, and that sensors should be deployed on systems themselves in addition to the network. He remarked that more needs to be done and called for terrain shaping, in which the cyber landscape is altered strategically to disadvantage attackers, not just to confuse them. Laddaga closed with a reminder that the need for AI to protect systems will increase as the cost of damage will move from data and machine damage to loss of human life.

The afternoon session of the workshop featured two interesting talks related to AI and its impact on cyber security. The first presented a call for a unified cyber ontology to help with information sharing. Such a mechanism represents concepts to share information automatically and could support deeper analytical investigation. The second talk presented a new mod-

el for effective cyber defense based upon active perception — deploying new sensors to understand the environment in the wake of a cyber event — to test hypotheses and reduce uncertainty prior to alerting an analyst.

The workshop concluded with a panel discussion on the subject AI and Cyber Operations: Challenges of Community Acceptance, including panelists Richard Lippmann (MIT Lincoln Laboratory), Vern Paxson (University of California at Berkeley), Benjamin Rubinstein (University of Melbourne), and Milind Tambe (University of Southern California). Panelists discussed the need to leverage natural language processing, game theory, and other AI capabilities as a toolbox within the specified domain of cyber security. The biggest challenge identified was that cyber is a rapidly changing domain, and it is extremely difficult to model the nonstationarity. Paxson stressed that his views on the limitations of machine learning in cyber security were specific to intrusion detection, which suffers from these challenges. Final conclusions from the panel focused on the need to leverage AI as a force multiplier and decision support tool for cyber defenders, focusing on human-in-the-loop rather than full automation.

This was the first AI for Cyber Security workshop. It was cochaired by David Martinez, William Streilein, Kevin M. Carter, and Arunesh Sinha, who also wrote this report. The papers of the workshop were published as technical report WS-16-03 in *The Workshops of the Thirtieth AAI Conference on Artificial Intelligence* (Palo Alto, CA: AAI Press).

Artificial Intelligence for Smart Grids and Smart Buildings

The proliferation of intelligent devices and the availability of electric monitoring facilities, broadband communication networks, computational intelligence, and customer-driven electricity storage and generation capabilities have posed the foundations for the next generation of power grids and buildings: smart grids and smart buildings. Three key aspects distinguish the

smart grid from the more traditional electric grid: (1) producers and consumers have access to information (for example, production costs, customers' electricity needs, time distribution of demands); (2) continuous access to information and communication is possible (for example, producers and consumers can negotiate prices); and (3) energy can be produced not only by power plants, but also by customers (for example, through renewable sources, which can be intermittent) and stored for later use (or redistributed through the electric grid). In general terms, a smart grid enables the distributed generation and two-directional flow of electricity, within an integrated system.

Smart buildings form an important component of the smart grid, where technology enables buildings to provide common services (for example, illumination, thermal comfort, air quality, sanitation) in a sustainable fashion and at low environmental impact. AI plays a key role in the smart grid and in smart buildings; the infrastructure provides information to support automated decision making on how to autonomously adapt production and consumption of energy, optimize costs, waste, and environmental impact, and ensure safe, secure, and efficient operation.

The goal of this workshop was to bring together researchers and practitioners from different areas of AI, to explore both established and novel applications of AI techniques to solve problems related to the design, implementation, and deployment of smart grids and smart buildings.

The workshop was well attended and conducive of extensive discussions and interactions between speakers and the audience. The workshop was opened by a team-delivered invited presentation by Mario Berges (Carnegie Mellon University) and Henning Lange (Aalto University), exploring the role of AI in solving problems related to electricity disaggregation — that is, the problem of providing estimates of the consumption of individual electrical appliances in a building from measurements of voltage and/or current at select locations in the facility. This initial presentation was followed by a related con-

tributed presentation, by the team from the University of Minnesota (Mark Valovage and Maria Gini), that focused on the use of label-correction techniques and prioritization methods in enhancing classification of individual appliances in a household during electricity disaggregation.

After a short break, the workshop continued with three contributed talks focused on the use of planning and optimization techniques in energy management, ranging from the use of planning with uncertainty to handle electric vehicles, to the use of MDP techniques in building agents for the Power TAC competition, to the exploration of how distributed constraint optimization can be adapted to handle the dynamicity of energy delivery models.

The afternoon session placed emphasis on the role of machine learning and classification in handling smart grids and smart buildings, with particular emphasis on reduced consumption prediction (in the presence of demand response events), prediction of adverse events and modeling of energy consumption in wireless sensor networks. The final session of the workshop explored personal preferences and individual comfort, with applications to support smart house buying (accounting for energy costs) and prediction of thermal comfort in smart buildings.

The workshop was cochaired by Enrico Pontelli and Tran Cao Son (New Mexico State University), Alex Rogers (University of Oxford), and Sylvie Thiebaux (NICTA and Australian National University), who also coauthored this report. The papers of the workshop were published as technical report WS-16-04 in *The Workshops of the Thirtieth AAI Conference on Artificial Intelligence* (Palo Alto, CA: AAI Press).

Beyond NP

A new computational paradigm has emerged in computer science over the past few decades, which is exemplified by the use of SAT solvers to tackle problems in the complexity class NP. According to this paradigm, a significant research and engineering invest-

ment is made toward developing highly efficient solvers for a prototypical problem (for example, SAT), that is representative of a broader class of problems (for example, NP). The cost of this investment is then amortized as these solvers are applied to a broader class of problems through reductions (in contrast to developing dedicated algorithms for each encountered problem). SAT solvers, for example, are now routinely used to solve problems in many domains, including diagnosis, planning, and software and hardware verification.

Motivated by the success of this computational paradigm, both in theory and practice, the goal of this workshop was to consolidate and promote research that advances this paradigm more broadly, while focusing on solvers that reach beyond NP. The workshop brought together researchers that work on a variety of problems in beyond NP complexity classes and corresponding solvers, including propositional and first-order probabilistic reasoning, planning, quantified Boolean formulas, function and optimization problems such as Max-SAT, knowledge compilation, and model counting. The workshop included six invited talks that provided a perspective on a wide spectrum of beyond NP solvers and problems. It also included 15 papers on these subjects, ranging from ones that included new technical contributions to ones that provided surveys and perspectives on the state of the art.

A number of themes emerged throughout the workshop and its talks. One theme concerned the status of competitions and evaluations for beyond NP solvers, the existence of standard formats and benchmarks for beyond NP problems, and the impact this may have on future progress. Another theme was the extent to which solvers for a particular problem have converged on a standard approach just like SAT solvers have converged on a standard approach for solving satisfiability. A major theme of the workshop related to the two facets of the PSPACE complexity class, which are exemplified by the two prototypical problems corresponding to quantified Boolean formulas and planning.

The workshop participants discussed

some next steps for advancing the reduce-then-solve computational paradigm targeted by the workshop. This included a categorization and explanation of prototypical problems that are complete for various beyond NP complexity classes, including the PP, NP^{PP}, and PP^{PP} complexity classes, which include key problems that arise in probabilistic reasoning. There was recognition that more community awareness is needed about beyond NP complexity classes and that existing and well-researched problems fall into these classes. It was also noted that the practice of reducing problems, which is now prevalent for problems in NP, can benefit from some tutorials and illustrative examples relating to beyond NP problems, especially ones that are PSPACE-complete. A third discussion point related to establishing standard formats and benchmarks for the more recent beyond NP solvers, including those for first-order model counting and knowledge compilation. The suggestion was to implement these next steps by augmenting the BeyondNP.org community website with additional material to serve these objectives.

Adnan Darwiche, Joao Marques-Silva, and Pierre Marquis served as cochairs of this workshop and wrote this report. The papers of the workshop were published as technical report WS-16-05 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press). The slides for invited talks are posted at the workshop website (beyondnp.org/workshop16).

Computer Poker and Imperfect Information Games

Recent years have brought 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 achieved human-level performance even 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 based only on a formal description of a game's 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 aspect and focusing on various domains in imperfect information games.

The workshop program was composed of 12 technical paper presentations, announcement of the results of the Annual Computer Poker Competition (ACPC), and two discussion sessions. The presentations covered topics from computing Nash and Stackelberg equilibria in large extensive form games and automated creation of game abstractions through opponent modeling and exploitation, heuristic search methods for imperfect information games, and taxonomy of different subclasses of imperfect information games, to applications of game theory in urban crime prevention and wildlife protection.

The winner of the total bankroll part of the 2016 computer poker competition was a program by Noam Brown and Tuomas Sandholm from Carnegie Mellon University. The winner of the bankroll instant run-off was implemented by an independent researcher, Eric Jackson.

The first discussion was on the future of the workshop. The attendees agreed that they are most interested in running a full-scale, six player No-Limit Texas Hold'em competition and that the total bankroll track of the competition should be modified to better motivate creating adaptive agents. Furthermore, the present competitors agreed that they do not mind using complicated variance reduction techniques in the evaluation of the algorithms and that it is important to attract more competitors to the competition. Interesting options to do that would be to include a track with a very small limit for the size of the agent, or a track that

cannot be entered by anyone who has already participated in multiple past competitions.

The topic of the second discussion was how to make the communities of researchers in poker, search in games, general game playing, and security games collaborate more efficiently. Everybody agreed that the problems the communities try to solve have many similar aspects and that they would benefit from closer interaction. The main drawback is that all the communities are quite productive and use slightly different terminology; therefore, it is difficult to follow the latest developments in the related fields. A proposed solution is to repeat events similar to this workshop that bring the communities together, and to present intensive advanced tutorials on the individual subfields. Unlike traditional tutorials, these could assume basic knowledge of key concepts of game theory and focus on the most important recent results and key challenges that the subfields face.

Viliam Lisy, Michael Thielscher, and Thanh Nguyen served as cochairs of this workshop. This report was written by Viliam Lisy. The papers of the workshop were published as technical report WS-16-06 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Declarative Learning Based Programming

Declarative Learning Based Programming (DeLBP) aims at facilitating and simplifying the design and development of intelligent real-world applications that use machine learning and reasoning by addressing the following commonly observed challenges: interacting with messy, naturally occurring data; specifying the requirements of the application at a high abstraction level; dealing with uncertainty in data and knowledge in various layers of the application program; using representations that support flexible relational feature engineering; using representations that support flexible reasoning and structure learning; integrating a range of learning and inference algorithms; and finally addressing these

issues in one unified programming environment.

Conventional programming languages offer no help to application programmers that attempt to design and develop applications that make use of real-world data and reason about it in a way that involves learning interdependent concepts from data, incorporating existing models, and reasoning about existing and trained models and their parametrization. Over the last few years the research community has tried to address these problems from multiple perspectives, most notably various approaches based on probabilistic programming, logical programming, and the integrated paradigms. The goal of this workshop was to present and discuss the current lines of research in these directions and the ways various challenges have been addressed. We attempted to motivate the need for further research toward a unified framework in this area, building on some of the key existing paradigms, including those of probabilistic programming (PP), probabilistic logical programming (PLP), database management systems (DBMS), and statistical relational learning (SRL), and place these ideas in the context of learning based programming.

This workshop brought together researchers from the areas of probabilistic logical programming, statistical relational learning, lifted inference, constraint programming, probabilistic programming, and data mining. The workshop started with an introductory talk given by Dan Roth that highlighted different perspectives and avenues for this research that lead to learning based programming and included two invited talks on these topics. Kristian Kersting talked on the democratization of optimization, focusing on declarative programming for lifted inference and efficient optimization in relational domains, exploiting the rich logical structure underlying many AI and data mining problems. Lise Getoor's talk introduced probabilistic soft logic (PSL), a declarative probabilistic programming language that can capture rich structure and scales well, and emphasized the PSL's mathematical framework, hinge-loss Markov random fields (HL-MRFs).

We ended the workshop with a panel and discussed the difficulties of developing and popularizing the use of these kinds of languages, the type of audience that we may want to target, interesting applications that can be addressed, and the ways different research communities can collaborate to make progress in this research direction.

Parisa Kordjamshidi served as the chair of this workshop. This report was written by Parisa Kordjamshidi and Dan Roth. The papers of the workshop were published as technical report WS-16-06 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Expanding the Boundaries of Health Informatics Using AI

The 20th century laid a foundation of evidence-based medicine that relied on populations and large groups of patients to derive generalized results and observations that were applied to (mostly passive) patients. Yet, the 21st century is shaping up as a time where the patient and personalized health data drive health care innovation and delivery. The availability of this vast amount of personalized data allows for care tailored to a specific patient, an approach coined *personalized medicine*. Moreover, the availability of this data allows for the constant monitoring and discovery of deviations from patient-specific averages (possibly different from population-based averages). These deviations may signal developing problems, and their early detection allows for more effective treatment leading to proactive medicine. Finally, patients are no longer passive recipients of (personalized) treatments and therapies, but they actively participate as a decision maker in their development, customization, and application. This shift has led to the emergence of participatory medicine.

To tackle issues that arise in proactive, personalized, and participatory medicine information technology will need to evolve to improve communication, collaboration, and teamwork among patients, their families, health-

care communities, and care teams involving practitioners from different fields and specialties. All of these changes require novel solutions and the AI community is well positioned to provide both theoretical- and application-based methods and frameworks. The goal of this workshop is to focus on creating and refining AI-based approaches that (1) process personalized data, (2) help patients (and families) participate in the care process, (3) improve patient participation, (4) help physicians utilize this participation in order to provide high quality and efficient personalized care, and (5) connect patients with information beyond those available within their care setting. The extraction, representation, and sharing of health data, patient preference elicitation, personalization of “generic” therapy plans, adaptation to care environments and available health expertise, and making medical information accessible to patients are some of the relevant problems in need of AI-based solutions.

This year’s workshop built on the very successful AAAI-13 workshop and AAAI 2014 fall symposium on the same topic. The workshop received a large number of submissions that were divided into two main themes. Learning and prediction focused on novel methods for mining semantics from patterns over electrocardiogram data, an adaptive ensemble learning approach for personalized diagnosis, predictive analytics using smartphone sensors for depressive episodes, and a learning-based approach to predicting the 30-day risk and cost of hospital readmissions. The information integration theme saw papers describing how to combine multiple concurrent physiological streams to assess a patient’s condition, and an approach to simultaneous influencing and mapping social networks to improve interventions for the homeless community in Los Angeles.

In addition to these themes, two invited speakers provided crucial insights into and directions for health informatics research. The first invited talk by Niels Peek (University of Manchester) was titled “Analytical Challenges for Smarter Health Systems.” Peek described the current opportunity

to create smarter, “learning” health systems by utilizing the information infrastructure provided by electronic health record (EHR) systems. Examples included using the data that is collected through this infrastructure to develop predictive models for risk stratification and to compare the effectiveness of different treatments in real-world populations. He also positioned EHR systems to be used as a delivery platform to give feedback and advice to clinicians at the point of care. Peek also challenged the audience to develop new tools to make sense of these data that reach beyond traditional analytical concepts and to improve clinical computerized decision support systems that have mostly led to alert fatigue rather than improving care.

The second talk, given by John H. Holmes (University of Pennsylvania), was titled “Data Driven Clinical Research: If Only It Were So Simple.” Holmes presented both the opportunities and challenges posed by the availability of ever-increasing amounts of highly heterogeneous clinical data for the data scientist and clinical researcher. The opportunity for enhanced clinical research is manifested in the expanding data and information ecosystem. The challenges are more subtly detected, but present nonetheless. Merging heterogeneous data into an analyzable whole, understanding the clinical context of an image or waveform without their semantic integration with clinical observation data, ecologic fallacy, and data quality were some of the challenges discussed.

Martin Michalowski served as the workshop chair and authored this report. Szymon Wilk and Jay M. Tenenbaum and served as cochairs. The papers of the workshop were published as technical report WS-16-08 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Incentives and Trust in Electronic Communities

The area of trust and reputation modeling has experienced rapid growth in the past decade. With the growing prevalence of social interaction

through electronic means, trust, reputation, and privacy, become considerably important. Many computational and theoretical models of trust and reputation mechanisms have been recently developed well-suited for variety of domains such as e-commerce, social network, blogs, ad hoc networks, and others. They present trust as a multifaceted concept that operates at many levels and plays important roles in ensuring reliable interactions. Although trust-enabled systems allow people to act under uncertainty and mitigate the risk of negative consequences, still socio-technical attacks often succeed by exploiting loopholes in the design of trust and security policies. Besides, the diversity of participants in the continuously growing electronic communities encourages cheating and opportunistic behaviors as it is more difficult in such environments to detect and punish fraudulent users. Many techniques have been developed to discourage deception and fraud in e-communities and stabilize trust between participants. These techniques are designed to promote trusting relationships, honesty behaviors, and create incentive for participants to contribute truthful opinions.

Trust and incentive have bidirectional relationships. As trustworthiness measures are used as part of incentive mechanisms to promote honesty in electronic communities, incentive mechanisms motivate participants to contribute their truthful opinions, which are useful for trust modeling. Hence, trust and reputation systems should not only provide a means to detect and prevent malicious activities but also design a mechanism to discourage dishonesty attitudes among participants.

The primary objective of this workshop is to bring together researchers in both the area of game theory for designing incentive mechanisms and the area of trust and reputation modeling, toward the design of more effective trust, reputation, and incentive mechanisms for creating safe electronic communities.

This year, our workshop accepted seven papers. There were also two invited talks, one given by Jiliang Tang (Yahoo Labs, USA) and another by Yan

Wang (Macquarie University, Australia). The workshop participants discussed the importance of designing an incentive-enabled system in electronic communities to foster honest participation among community members and how the workshop was useful in bringing together researchers from different fields such as game theory, user modeling, and trust modeling and encouraging their contributions in creating safe electronic communities.

Jie Zhang, Zeinab Noorian, and Stephen Marsh served as cochairs of this workshop. This report was written by Jie Zhang and Zeinab Noorian. The papers of the workshop were published as technical report WS-16-09 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Knowledge Extraction from Text

Text understanding is an old, yet-unsolved, AI problem consisting of a number of nontrivial steps. The critical step in solving the understanding problem is knowledge acquisition from text, that is, a transition from a nonformalized text, explicitly or implicitly, into a formalized actionable language (that is, capable of supporting automated reasoning). Other steps in the text-understanding pipeline include linguistic processing, reasoning, text generation, search, question answering, and others. These are more or less solved to a degree that would support composition of a text understanding service. However, we know that knowledge acquisition, the key bottleneck, can be done by humans, even though automation of the process is still out of reach in its full breadth.

After failed attempts in the past (due to a lack of both theoretical and technological prerequisites), in recent years the interest in text understanding and knowledge acquisition from text has been growing. There are numerous AI research groups studying various aspects of the problem in the areas of computational linguistics, machine learning, probabilistic and logical reasoning, and semantic web. The commonality among all the newer approaches is the use of recent

advances in machine learning to deal with representational change on the level of words, sentences, concepts, and so on.

The workshop brought together researchers from a variety of different approaches for extracting knowledge from text, in addition to researchers in fields that provide empirical or theoretical foundations to the main topic of the workshop. The oral presentations (including the keynote given by Peter Clark, from AI2) revealed a set of interesting and innovative ideas, as well as topics that should help to guide the future of the research community. The workshop participants discussed how knowledge extraction approaches should go beyond named entity recognition (NER) and relation-extraction tasks. The scientific debates focused on new types of knowledge that can be extracted from text, such as relations among events instead of relations among entities, as well as what the community has learned about extracting knowledge from specific target domain texts, such as mars literature. The challenge of extracting knowledge from texts written in different languages (other than English), such as Chinese and Portuguese, was also addressed, focusing on different domains such as never-ending learning to read the web, commonsense relations extraction, as well as automatic email answering agents. The workshop also presented discussions on how parallel computing can help in the scaling-up of topic model algorithms.

Blaz Fortuna, Marko Grobelnik, Estevam Hruschka, and Michael Witbrock served as cochairs of this workshop and wrote this report. The papers of the workshop were published as technical report WS-16-10 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Multiagent Interaction Without Prior Coordination

Interaction between agents is the defining attribute of multiagent systems, encompassing problems such as planning in a decentralized setting, learn-

ing other agent models, composing teams with high task performance, and selected resource-bounded communication and coordination. While there is significant variety in methodologies used to solve such problems, many of these methods depend on some form of prior coordination. For example, learning algorithms may assume that all agents share a common learning method and prior beliefs, distributed optimization methods may assume specific structural constraints regarding the partition of state space and cost/rewards, and symbolic methods often make strong assumptions regarding norms and protocols. However, such assumptions are easily violated in realistic problems. Thus, there is a need for new models and algorithms that specifically address the case of ad hoc interactions.

The purpose of this workshop was to discuss the role of such predefined knowledge and coordination in multiagent systems, and to provide a venue for research on novel models and algorithms that specifically address multiagent interaction without prior coordination (MIPC). There were a total of six accepted papers, addressing diverse topics such as ad hoc coalitions in human-robot societies, identifying and tracking nonstationary opponents, and policy communication for coordination with unknown teammates. A continuing trend in several of the workshop papers was the use of beliefs over a set of hypothesized policy types. In addition to the paper presentations, there were invited talks by Michael Bowling from the University of Alberta on "Adventures in Implicit Agent Modeling," and Gal Kaminka from Bar-Ilan University on "Teams, Swarms, Crowds, and Collectives: Special Cases?"

This was the third meeting in this workshop series. The community has grown steadily since the first workshop took place in 2014, reflecting a growing awareness of issues relating to prior coordination in multiagent systems. The workshop organizers intend to continue the workshop series. The workshop was chaired by Stefano Albrecht, Katie Genter, and Somchaya Liemhetcharat. The chairs would like to thank the workshop participants,

the invited speakers, the program and advisory committee, and the AAAI staff for making the workshop a success.

This report was written by Stefano Albrecht. The papers of the workshop were published as The papers of the workshop were published as technical report WS-16-12 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Planning for Hybrid Systems

The purpose of this workshop was to explore and promote new approaches to planning with hybrid models. Hybrid systems are systems with both continuous control variables and discrete logical modes. Many interesting real problems are indeed hybrid systems, including oil refinery management, mission planning for autonomous vehicles, supply management and disaster recovery, and applications in control of smart cities. Planning in these domains requires rich models to capture the interaction between discrete and continuous change and methods for reasoning with temporal, spatial, and continuous constraints.

This workshop follows on two previous highly successful hybrid planning events held at the International Conference on Automated Planning and Scheduling (ICAPS). ICAPS 2012 ran a special track on continuous planning, attracting a range of excellent talks on the topic of integrated planning and control. This track brought together planning experts, roboticists, experts in model-based reasoning, experts in runtime verification, and control engineers. This continued as a well-attended workshop at ICAPS 2013. The workshop was opened up to the wider AI community (and applications of interest to this broader community) by the workshop at AAAI-16, which was attended by researchers from planning, robotics, machine learning, hybrid system control and verification, and model-based reasoning and led to a fascinating range of papers and talks. Key topical areas for paper submission included architectures for hybrid systems, planning with SMT (satisfiability modulo theories), temporal logic in

hybrid systems, PDDL+ planning, and a fascinating range of applications spanning quantum computing for Mars Lander activity scheduling, urban traffic signal control, autonomous marine vehicles, airport surface operations, and robotics. In short, the exceptional attendance and energetic discussion of this workshop demonstrate the growing influence and applications relevance of hybrid systems planning in AI along with its appeal to a diverse range of researchers and research fields.

This workshop was organized and cochaired by Daniele Magazzeni (King's College London, UK), Scott Sanner (Oregon State University), and Sylvie Thiebaut (Australian National University and NICTA/Data61) and held on February 13, 2016, in Phoenix. Daniele Magazzeni, Scott Sanner, and Sylvie Thiebaut also wrote this report. The papers of the workshop were published as technical report WS-16-12 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Scholarly Big Data: AI Perspectives, Challenges, and Ideas

Academics and researchers worldwide continue to produce large numbers of scholarly documents including papers, books, technical reports, and associated data such as tutorials, proposals, and course materials. For example, PubMed has over 20 million documents, 10 million unique names, and 70 million name mentions. Google Scholar has many millions more, it is believed. Understanding how at scale research topics emerge, evolve, or disappear, what is a good measure of quality of published works, what are the most promising areas of research, how authors connect and influence each other, who are the experts in a field, and who funds a particular research topic are some of the major foci of the rapidly emerging field of scholarly big data.

The primary goals and objectives of the workshop were to promote both theoretical results and practical applications for scholarly big data, and address challenges faced by today's researchers, decision makers, and fund-

ing agencies as well as well-known technological companies such as Microsoft and Google, repositories, and publishers such as Elsevier.

Papers presented at the workshop covered a variety of topics including the presentation of a test collection for citation recommendation; the design of evaluation data sets for document similarity models in large scholarly retrieval systems; an analysis of NIH funding patterns over time; data extraction from scientific charts and summary generation; an approach to extend research footprints using disparate sources; improving public access to nonopen biomedical literature; topic-level academic influence of scientific literature; lineage encoding in scholarly articles; and improving discoverability of research papers by augmenting their titles with more terms.

The workshop also included two invited talks and an introduction to the CiteSeerX digital library. The first talk, given by Douglas Downey (Northwestern University), focused on mining topics and key phrases from scientific documents and the application of keyphrase extraction within the semantic scholar scientific search engine. The second talk, given by Alex Wade (Microsoft Research Redmond), focused on academic knowledge: new research opportunities with the Microsoft Academic Graph (MAG), which is a freely available data set that includes information about academic publications and citations, researchers, venues, and topics. MAG is a heterogeneous graph that can be used to study the influential nodes of various types, including authors, affiliations, and venues.

Cornelia Caragea, C. Lee Giles, Alex Wade, and Irwin King served as cochairs of this workshop. This report was written by Cornelia Caragea and C. Lee Giles. The papers of the workshop were published as technical report WS-16-13 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Symbiotic Cognitive Systems

In his 1960 article "Man-Machine Symbiosis," J. C. R. Licklider predicted a time when "the main intellectual

advances will be made by men and computers working together in intimate association.” While much of the emphasis within the AI community over the ensuing half century was placed upon tools for automation such as speech recognition or surpassing humans at challenging intellectual tasks such as chess or *Jeopardy!* the last few years have witnessed a resurgent interest in symbiotic cognitive systems: collectives of humans and intelligent agents that collaborate to accomplish cognitive tasks better than either can alone.

The objective of this workshop was to synthesize a new vision and research agenda for symbiotic cognitive systems and to try to establish a community that might have a continued existence at future AI workshops and conferences. For that purpose, the workshop brought together researchers working on technologies, architectures, applications, and visions of symbiotic cognitive computing in application domains that spanned robotics, cognitive environments, and cognitive objects. Presentations at the workshop fell into three main categories, with one session devoted to each, with ample time devoted at the end of each session to discussing, summarizing and abstracting the presented material.

In the session on patterns of symbiotic systems, presenters discussed general issues that apply broadly across many different instantiations of symbiotic cognitive systems. Several discussed aspects of knowledge representation, including its role in enabling distributed cognition, in semantically mapping elements of the physical world to that representation, and in teaching agents about relationships that need to be included in that representation. One insight that emerged during the discussion was the idea of confusion or uncertainty. The ability of a system to sense its own uncertainty is necessary in order to generate behaviors that are likely to improve the system’s representation of the world or the humans and agents in the environment. The role that attention can play in reorienting the priorities of a cognitive system was discussed, as was the relationship of attention to representation.

The session on symbiotic interactions focused on aspects of interactions between machines and humans or the world that they inhabit. One thread of discussion concerned verbal communication of needs, intents and plans. Methods of collecting verbal communication methods included crowdsourcing to collect a variety of language, profiling human users using social norms and elicitation dialogs, and virtual reality. A second thread of discussion concerned nonverbal communication of information pertaining to the physiological, mental, or emotional state of humans, possibly using cognitive objects (such as chairs) as sensors, while a third concerned systems that could be taught to interpret and execute commands that involved recognizing and manipulating objects in the physical world.

The third session was on learning about and from humans. On the topic of learning about humans, work was presented and discussed about inferring fatigue or other physiological attributes from eye tracking, and learning models of human reaction to advertising by observing behavior and analyzing facial expressions. The use of human demonstrations, situated dialog, text understanding, virtual words, and storytelling were all discussed as methods to learn from humans. The third session was concluded by an interactive poster session where authors could present and discuss their work in further detail.

The workshop concluded with a lively discussion among all participants that focused on questions such as “What are the essential common characteristics and themes of symbiotic cognitive systems?” “What distinguishes them from other interactive systems?” and “Is it worthwhile to create a symbiotic community, and if so how should we go about it?” It is fair to say that, while an absolute consensus was not reached regarding the first two questions, there was great enthusiasm for continuing the discussion at another venue in the near future, perhaps at AAMAS, IJCAI, or another AI conference.

This report was written by Jeffrey O. Kephart and Stephanie Rosenthal. Jeffrey O. Kephart, Stephanie Rosenthal,

and Manuela Veloso served as cochairs of the workshop. The papers of the workshop were published as technical report WS-16-14 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

World Wide Web and Population 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.

This workshop follows the success of previous AAAI workshops on the same topic held in 2014 in Quebec, Canada, and in 2015 in Austin, Texas, USA. This workshop brought together computer scientists, biomedical and health informaticians, researchers, students, industry professionals, and representatives of national and international public health agencies. Participants were interested in the theory and practice of computational models of web-based public health intelligence. 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, and online social media analytics. From an application perspective, presentations addressed topics in epidemiology, environmental and public health informatics, disease surveillance, health behavior monitoring, and disaster management.

One of the main themes of this workshop was the exploration and monitoring of online social media (for

example, twitter) to analyze behavioral patterns. Models of behavior were used to enhance forecasting, guide decision making, enable situational awareness, and inform response strategies. The workshop also included three invited talks. Sudha Ram (professor and director of INSITE Center for Business Intelligence and Analytics, University of Arizona) gave a presentation on using big data for predictive analytics in population and personalized health care. She demonstrated examples from developing predictive models using streaming sensor and social media data sets combined with health-care records. Soon Ae Chun (professor and director of the Information Systems Informatics program at the City University of New York, College of Staten Island) described values and risks of using patient-generated social health data in health care. Damon Centola (associate professor of communication at the Annenberg School for Communication at the University of Pennsylvania) presented his findings from a series of novel experiments designed to study the dynamics of behavioral diffusion in large social networks. His results showed a striking effect of network topology on the diffusion of health behavior, contrary to the expectations of classical network theory.

To promote open debate and exchange of opinion among participants, the workshop was concluded with a panel discussion moderated by Arash Shaban-Nejad and including Sudha Ram, Soon Ae Chun, and Damon Centola. The major theme of the panel was to discuss the future of online surveillance and interventions for changing individual health behaviors.

Arash Shaban-Nejad, David L. Buckeridge, Byron C. Wallace, and John S. Brownstein served as co-chairs of this workshop and wrote this report. The papers of the workshop were published as technical report WS-16-15 in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence* (Palo Alto, CA: AAAI Press).

Stefano Albrecht is a postdoctoral fellow at the University of Texas at Austin.

Bruno Bouchard is assistant professor at

Université du Québec à Chicoutimi .

John S. Brownstein is a researcher for Boston Children's Hospital and a professor at Harvard University.

David L. Buckeridge, MD, is an associate professor at McGill University.

Cornelia Caragea is an assistant professor at University of North Texas.

Kevin M. Carter is an associate group leader at MIT Lincoln Laboratory.

Adnan Darwiche is a professor at the University of California, Los Angeles.

Blaz Fortuna is a researcher consultant for Bloomberg L.P. and Jozef Stefan Institute.

Yannick Francillette is from Université du Québec à Chicoutimi

Sébastien Gaboury is assistant professor at Université du Québec à Chicoutimi.

C. Lee Giles is a professor at the Pennsylvania State University.

Marko Grobelnik is a research team leader at Jozef Stefan Institute.

Estevam R. Hruschka Jr. is an associate professor at the Federal University of São Carlos.

Jeffrey O. Kephart is a member of the research staff at IBM Thomas J. Watson Research Center.

Parisa Kordjamshidi is a postdoctoral researcher at the University of Illinois at Urbana-Champaign.

Viliam Lisy is a postdoctoral fellow at the Department of Computing Science, University of Alberta.

Daniele Magazzeni is a lecturer in computer science at King's College London.

Joao Marques-Silva is a professor at the University of Lisbon.

Pierre Marquis is a professor at Université d'Artois.

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

Martin Michalowski is a senior principal research scientist at Adventium Labs.

Zeinab Noorian is a postdoctoral research fellow at Ryerson University.

Enrico Pontelli is a professor at New Mexico State University.

Alex Rogers is a professor at the University of Oxford.

Stephanie Rosenthal is a research scientist at Carnegie Mellon University.

Dan Roth is a professor at the University of Illinois at Urbana-Champaign.

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

Arash Shaban-Nejad is affiliated with the School of Public Health, University of California, Berkeley.

Arunesh Sinha is a postdoctoral researcher at the University of Southern California.

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

William Streilein is a Group Leader at MIT Lincoln Laboratory.

Sylvie Thiebaut is a professor at the Australian National University and a research leader at the National ICT Australia.

Byron C. Wallace is an assistant professor at the University of Texas at Austin.

Toby Walsh is a professor at the University of New South Wales and a group leader at Data61.

Michael Witbrock is a consulting computer scientist at Lucid AI.

Jie Zhang is an associate professor at Nanyang Technological University.