Reports of the AAAI 2014 Conference Workshops

Stefano V. Albrecht, André M. S. Barreto, Darius Braziunas, David L. Buckeridge, MD, Heriberto Cuayáhuitl, Nina Dethlefs, Markus Endres, Amir-massoud Farahmand, Mark Fox, Lutz Frommberger, Sam Ganzfried, Sébastien Guillet, Yolanda Gil, Lawrence E. Hunter, Arnav Jhala, Kristian Kersting, George Konidaris, Freddy Lecue, Sheila McIlraith, Sriraam Natarajan, Zeinab Noorian, David Poole, Rémi Ronfard, Alessandro Saffiotti, Arash Shaban-Nejad, Biplav Srivastava, Gerald Tesauro, Rosario Uceda-Sosa, Guy Van den Broeck, Martijn van Otterlo, Byron C. Wallace, Paul Weng, Jenna Wiens, Jie Zhang

■ *The AAAI-14 Workshop program was held Sunday and Monday, July 27–28, 2014, at the Québec City Convention Centre in Québec, Canada. The AAAI-14 workshop program included 15 workshops covering a wide range of topics in artificial intelligence. The titles of the workshops were Artificial Intelligence and Robotics; Artificial Intelligence Applied to Assistive Technologies and Smart Environments; Cognitive Computing for Augmented Human Intelligence; Computer Poker and Imperfect Information; Discovery Informatics; Incentives and Trust in Electronic Communities; Intelligent Cinematography and Editing; Machine Learning for Interactive Systems: Bridging the Gap Between Perception, Action, and Communication; Modern Artificial Intelligence for Health Analytics; Multiagent Interaction Without Prior Coordination; Multidisciplinary Workshop on Advances in Preference Handling; Semantic Cities — Beyond Open Data to Models, Standards, and Reasoning; Sequential Decision Making with Big Data; Statistical Relational AI; and the World Wide Web and Public Health Intelligence. This article presents short summaries of those events.*

Artificial Intelligence and Robotics

The fields of AI and robotics were strongly connected in the early days of AI, but have since diverged. One of the early goals of AI consisted of building complete intelligent robots. However, such a goal was shown to be quite challenging, and AI and robotics researchers therefore isolated its many facets and focused on making progress on each facet separately.

Both AI and robotics have matured enormously since those days, and today there is a growing interest in getting the two fields back together. Many in robotics believe that the next quantum leap in that field will come by adding capabilities that lie at the core of AI research like task planning, knowledge representation, learning, and natural interaction. Symmetrically, AI researchers show interest in embedding their techniques in platforms that can perceive, reason, and act in real, dynamic physical environments.

Despite this mutual interest, there is a shortage of mainstream venues where the AI and robotics communities can meet. Researchers in either field are often unaware of the methodology, the successes, and the limitations in the other one. More generally, we do not have enough understanding of what are the fundamental issues in putting together an integrated intelligent robotic system. Examples of such systems exist, but they remain isolated points in a landscape whose overall structure and extent are still not clear. The objective of the Artificial Intelligence and Robotics workshop was to contribute to the creation of a joint community of AI and robotics researchers that can fill this gap.

The workshop gathered more than 30 participants, who came from different communities but felt the urge to combine the AI and robotics traditions. These included AI researchers interested in robotics, robotics researchers interested in AI, and researchers who are genuinely working in the middle ground.

The workshop was introduced by a general talk by the organizers about the challenges of opportunities of integrating AI and robotics. The 10 contributed papers touched some of the major aspects of the integration between AI and robotics: the link between sensor data and high-level descriptions of activities or of spatial relations; the use of attention to improve perception; and learning in vision and in control. Several papers addressed different facets of the problem of deliberating and acting in the real world. These included the use of planning to create reactive, collective control policies; plan execution and plan monitoring; and a study on autonomous goal generation. Some of the papers were complemented by a poster in order to provide more opportunities for discussion. The workshop also included a special "taking issue" session devoted to the discussion of papers that raised general and controversial issues, like how to evaluate empirically an intelligent robotic system.

This workshop was part of a series of initiatives whose aim is to bring the two areas of artificial intelligence and robotics closer together, fostering the creation of a joint community, of joint venues, and of a joint education program that can breed a new generation of researchers who can bridge the two disciplines. The last part of the workshop was devoted to a discussion about this community effort. A key tool for this effort is the AI and robotics wiki at ai-robotics.wikispaces.com, which includes a list of relevant scientific and educational initiatives.

This report was written by Alessandro Saffiotti and George Konidaris. Nick Hawes, George Konidaris, Alessandro Saffiotti, and Moritz Tenorth served as organizers of this workshop. A program committee of 26 people helped to review the submissions and create the program: their names are listed in the technical report. The papers presented at the workshop were published as AAAI Press Technical Report WS-14-01.

Artificial Intelligence Applied to Assistive Technologies and Smart Environments

The goal of the Artificial Intelligence Applied to Assistive Technologies and Smart Environments workshop was to expose the research area of assistive technologies to artificial intelligence specialists, and more especially to explore the practical applications of recent artificial intelligence techniques to create, deploy, and — above all — adapt assistive technologies and smart environments to enable users with physical or cognitive impairments — a group that frequently includes elderly people — to maintain their independence and remain in their homes longer.

The workshop brought together researchers from a variety of subfields of AI such as multiagent systems and decision support, in addition to researchers in fields providing empirical or theoretical foundations (including personalized assistance, driving assistance, activity prediction and recognition, and high-level control of autonomous systems). One major theme of papers presented at the workshop was ambient assisted living (AAL). The workshop also included one invited talk on this theme. A first talk given by Juan Carlos Augusto (Middlesex University, London, UK) focused on the importance of preferences and needs to increase the perceived usefulness of assistive technology applications from the final user side. His second talk described an AAL system and its use in the outdoors to enhance independent living. A talk given by Patrice C. Roy (Dalhousie University, Halifax, Canada) motivated the use of standard clinical guidelines to implement AAL systems.

Another major theme was the security of the user. The papers on this theme described systems that can prevent or adapt to user errors. The papers of Tom Croonenborghs (KU Leuven, Belgium) and Colombiano Kedowide (Télé-Université du Québec, Montréal, Canada) especially focused on the use of AI techniques to automatically detect incidents, while the paper of Sébastien Guillet (UQAC, Chicoutimi, Canada) outlined general principles to model and synthesize the control part of AI systems that can adapt a smart home to prevent potentially harmful situations.

The third theme of the workshop was the use of localization for activity recognition and assistance. Two papers presented by Kevin Bouchard (UQAC, Chicoutimi, Canada) took advantage of RFID to get precise position of equipped objects and use these positions within a time-series framework to predict activities, respectively. The paper of Sylvain Giroux (Université de Sherbrooke, Canada) proposed a multiagent approach to follow people and provide personalized assistance based on their localization.

The workshop participants discussed how the trend toward the development of new assistive technologies is growing to help people with disabilities

but still not ready to be adopted by them. Based on this observation, the participants shared the goal of pursuing efforts in the activity-recognition domain, as it constitutes the foundation of AAL systems, from a low-level point of view (interpreting sensor information as activities) to higher levels (implementing assistance and security measures from recognized activities). The participants also agreed that they would like to attend future workshops with the same focus as this one.

This report was written by Sébastien Guillet. Bruno Bouchard served as chair of this workshop. Sylvain Giroux, Abdenour Bouzouane, Alex Mihailidis, and Sébastien Guillet served as cochairs. The papers of this workshop were published as AAAI Press technical Report WS-14-02.

Cognitive Computing for Augmented Human Intelligence

Inspired by major advances in AI in the last few years, a new research topic called cognitive computing has emerged independently across many institutions, including universities (MIT, Georgia Tech, University of Calgary), government funding agencies (DARPA, European Union), and technology companies (IBM, Microsoft, Google, Facebook, and others). While there is not yet a standard definition of cognitive computing, the common theme appears to be combining automatic learning and knowledge extraction from massive-scale data with much more naturalistic interactions between humans and systems. The roles of such systems will differ significantly from traditional AI systems that autonomously perform some desired task (such as playing chess or theorem proving). Instead, these systems will be inherently engaging and collaborative with humans, much akin to the way humans work together in teams. Achieving this vision could enable much better decision making and productivity enhancements across many industries, including education, health care, and management of large distributed systems such as transportation networks and electric power grids.

This workshop brought together 25–30 AI researchers from around the world to discuss a panoply of research topics with potential relevance to the cognitive computing vision. We had a strong program of contributed papers, representing quite diverse aspects of cognitive computing. There were two papers with a primary focus on natural language processing, three papers that studied visual perception tasks, two papers that investigated knowledge representation and reasoning techniques, and two papers focused on modeling the cognitive states of human users. Some of the noteworthy contributions were as follows. The paper by Yuetan Lin (Zhejiang University) and others proposed combining feature learning and saliency map learning within a convolutional network architecture to automatically detect the most salient objects in complex images. The paper by Jonathan Dunn (Illinois Institute of Technology) and others presented a novel ensemble-based approach to identifying metaphorical language in arbitrary text. This approach can be applied across multiple languages, and perhaps surprisingly, does not require usage of any language-specific resources. The paper by Yu-Ting Li and Juan Wachs (Purdue University) developed methods to learn Bayesian network models of human attention, given observable cues such as gestures and body stance. These models were validated in laboratory studies of human task performance, and it was found that predicted improvements in attention across various interaction modalities were associated with improved spatio-navigational problem solving.

The strongest contributions to success of the workshop arguably came from a stellar lineup of invited speakers, with quite disparate perspectives. Yoshua Bengio (University of Montreal) provided a comprehensive overview of major research challenges in deep learning, such as approximate inference, disentangling factors of variation, reasoning, and one-shot learning of facts. Milind Tambe (USC) discussed the challenges of bounded-rationality game-theoretic reasoning in the context of real-world human security games. Bonnie John (CMU) gave a compelling talk on the importance of addressing major challenges from the user perspective (usability, trust, and others) in building cognitive computing systems. Pat Langley (University of Auckland) delivered a powerful critique of research studying a single aspect of cognition (for example, perception) in isolation, and strongly advocated the agenda of cognitive systems, which aims to build artifacts exhibiting the full range of human intelligence. Richard Socher (Stanford University) discussed three fascinating topics at the forefront of deep learning research: knowledge base completion using tensor networks, zero-shot learning through cross-modal transfer, and sentence embeddings for describing images.

With such diverging views among the participants and invited speakers, it is not surprising that very lively dialogue and debate ensued during the open discussion periods. One of the most extensively debated topics concerned whether research success in cognitive computing is more likely to result from focusing on integration of many distinct modules (learning, planning, knowledge representation, perception, and others), each representing a single component of human intelligence, versus the traditional path of studying and advancing technology for individual components in isolation. Apart from the purely technical merits of the two approaches, participants also noted significant "career risk" in attempting to do integrative work that is unlikely to be accepted in prominent mainstream venues with specific focus in, for example, machine learning or NLP or computer vision. Despite any contentious

debates, workshop participants reached an amicable conclusion that followup workshops will be worthwhile, in order to build research momentum in what could be an extremely significant topic in the next several years.

Gerald Tesauro was the author of this report. Biplav Srivastava, Aurelie Lozano, Janusz Marecki, Irina Rish, Ruslan Salakhudtinov, Gerald Tesauro, and Manuela Veloso served as cochairs of the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-14-03. Slides of the invited talks are available at the workshop website.[1]

Computer Poker and Imperfect Information

In recent years, poker has emerged as an important, visible challenge problem for the field of AI. Just as the development of world-class chess-playing programs was considered an important milestone in the development of intelligent computing, poker is increasingly being seen in the same way. Several important features differentiate poker from other games: the presence of imperfect information (due to hidden cards), stochastic events, and the desire to maximize utility instead of simply winning. Games of imperfect information typically require randomized strategies that "hide information" effectively. For these reasons and others, games of imperfect information require methods quite different from traditional games of perfect information like chess or Go.

The goal of this workshop was to bring together researchers studying theoretical and practical aspects of imperfect-information games to share current research and gather ideas about how to improve the state of the art and advance AI research in this area. The workshop was initiated at AAAI 2012 as the Computer Poker Symposium, after which it turned into the Workshop on Computer Poker and Imperfect Information for 2013 and 2014.

The workshop featured 10 oral presentations, followed by a poster session and a roundtable discussion. The presentations addressed a range of topics on games of imperfect information. Several of the talks proposed new equilibrium-finding algorithms—including variations of Monte Carlo tree search and Monte Carlo counterfactual regret minimization, and a new interior point algorithm. Algorithms for game abstraction were another popular topic. Noam Brown from Carnegie Mellon University presented a new approach for warm starting no-regret learning algorithms and showed how this approach can be used to optimize game parameters, such as bet sizes in no-limit poker. Eric Jackson presented a new game decomposition algorithm for approximately solving large imperfect information games by breaking the problem into subproblems that can be solved independently. Trevor Davis from Alberta presented work on a new way to measure the strength of a strategy that goes beyond worst-case exploitability. Neil Burch from Alberta presented the results from the poker competition.

The 2014 competition had 12 different agents in the two-player limit Texas hold'em competition, 16 agents in the two-player no-limit competition, and 6 agents in the three-player limit competition. Agents were submitted by a mixture of universities and hobbyists from more than 10 countries. Following a poster session in which 10 posters were presented, there was a roundable discussion that focused on plans for future poker competitions.

The workshop will be held again at AAAI-15 chaired by Sam Ganzfried. Given the short gap between AAAI-14 and AAAI-15, the consensus was that it may be too soon to have a full competition in 2015, although there was a discussion of running some events in a smaller-scale 2015 competition, followed by another full competition in 2016. There was also a discussion of possible introduction or removal of certain variants from the competition, as well as application domains outside of poker to which algorithms developed for poker agents can be applied.

This report was written by Sam Ganzfried. Sam Ganzfried and Eric Jackson served as cochairs of this workshop. The papers of the workshop were published as AAAI Technical Report WS-14-04.

Discovery Informatics

The premise of this workshop was that intelligent systems are key to future advancements of science. With increasingly more data available, more complex methods at their disposal, and a very rapidly growing literature, scientists need intelligent assistants that can take on significant tasks in the scientific endeavor. This workshop convened researchers and practitioners working toward this vision.

Several participants presented their work in bringing AI techniques to address particular science challenges. These include the use of reasoning and inference to support semantic data integration, the use of semantic networks to integrate diverse sources of scientific knowledge, and the use of causal similarity to detect regularities in independent data streams.

Some presentations discussed the more philosophical aspects of the use of intelligent systems in science. This included the role of computational creativity in the generation of scientific hypotheses, computational ideation to develop and revise models, and the potential of scientific domains as sources of grand challenges for AI.

Other presentations highlighted the importance of improving science practice in representing more explicitly the information that can eventually help intelligent systems better understand scientific data and papers. The topics discussed included improving the reuse and documentation of research data, creat-

ing annotations of scientific publications in terms of their claims and argumentation structure, and formalizing the links between experimental evidence and interpretation.

The workshop had three invited speakers that reflected on major challenges and directions for AI in science. Tom Dietterich, distinguished professor in computer science at Oregon State University, described recent research on discovering bird migration patterns where scientific hypotheses are represented as latent variables. Paul Cohen, program manager at the Defense Advanced Research Projects Agency, discussed DARPA's new Big Mechanism program aimed to infer causal explanatory models for complex systems from large and diverse data sets. Phil Bourne, the newly appointed associate director for data science at the National Institutes of Health, discussed potential opportunities for the discovery informatics community to contribute to the new initiatives at NIH to integrate information in an open science commons.

Participants discussed a broad range of scientific domains, including biology, medicine, geosciences, environmental sciences, and material sciences. The research in this area tends to be presented in scientific conferences in those fields, and this workshop presented a unique opportunity for cross-fertilization of ideas and approaches in this community.

Yolanda Gil and Larry Hunter served as cochairs of this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-14-05.

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 that are well-suited for a variety of domains such as e-commerce, social networks, 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 designing trust-enabled systems allows 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 and honesty behaviors and to 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 that 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 amongst 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. The workshop participants discussed the importance of designing an incentive-enabled system in electronic community environments 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 and Zeinab Noorian wrote this report. Jie Zhang, Zeinab Noorian, Stephen Marsh, and Christian Damsgaard Jensen served as cochairs of this symposium.

Intelligent Cinematography and Editing

Intelligent cinematography and editing (ICE) is an emergent research field whose object of study is the expressive use of virtual cameras, mise-en-scène, lighting, and editing (montage) techniques by computers. ICE is important both as a basic research problem (can moviemaking be broken down into computer programs? can computers be creative?) and for its applications to computational media (for example, automatic generation of cut-scenes in games), education (for example, tutorial generation), and moviemaking (such as for novice filmmakers).

The workshop featured Mark Riedl as the keynote speaker. Riedl directs the Entertainment Intelligence Lab at the Georgia Institute for Technology. In his keynote talk, Riedl proposed the "grand challenge" of having a computer automatically generate a feature-length movie, and receive 20 percent positive opinions from an audience rating system such as Rotten Tomatoes.

Consistent with the stages of the film production pipeline, Riedl divided this grand challenge into the tasks of generating the plot (cognitive modeling, natural language processing, knowledge representa-

tion), writing the script (natural language generation), and shooting the movie (computer vision and graphics). Riedl's work on automatic plot generation starts from the claim that creative writing can be represented as a problem-solving process. Thus, his work seeks to use classical plan representations to express events that take place in a narrative to satisfy authorial goals. Previous work has shown that partial order planning is suitable to the difficult task of planning character actions that realize the goals of the storyteller. Riedl emphasized how the storytelling domain is different from domains typically used in traditional planning research, by using the example of Shakespeare's *Hamlet,* where the death of the main characters is clearly not part of their own goals, but part of the author's goals. The topics of intelligent cinematography and editing were covered by describing the CamBot system, which takes a symbolically encoded movie script and turns it into a three-dimensional animation movie.

Interestingly, the most difficult part in meeting the grand challenge now appears to be not in plot generation or in script-based cinematography and editing, but in filling the gap between the computer-generated plot and the computer-readable script. This promises to become a fertile topic in the future for the AI and machine-learning communities.

Technical papers highlighted a wealth of recent trends in intelligent cinematography. Peter Carr presented a survey of recent work on intelligent editing for live sports events. Cameron Alston and Arnav Jhala showcased a gaze tracking library with sample games, including a three-dimensional first-person shooter with gaze control. Cunka Sanokho and Marc Christie described a system that learns rules of balance in shot composition from annotated movie clips. Christophe Lino, Rémi Ronfard, Quentin Galvane, and Michael Gleicher offered a discussion of the evaluation problem in computer-generated film editing. Anderson and Packard argued that computational systems still fail to capture the creativity of human directors in creating discontinuity in visual storytelling to elicit viewer emotions. Devon Bates and Arnav Jhala offered definitions for audio and video pace, and their role in automatically identifying important narrative transitions depicted in scene transitions.

The poster session focused on applications to training and education, including an automated system for lecture recording, a declarative camera control language in a hierarchical planning formalism for cinematic vocabulary, and an improved shot-classification algorithm to detect narrative transitions in movie clips.

Overall, the workshop was successful in opening interesting interdisciplinary research avenues for the AI community. Rémi Ronfard and Arnav Jhala wrote this report. Rémi Ronfard, Arnav Jhala, and Paolo Burelli served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-14-06.

Machine Learning for Interactive Systems: Bridging the Gap Between Perception, Action, and Communication

Intelligent systems or robots that interact with their environment by perceiving, acting, or communicating often face a challenge in how to bring these different concepts together. Applications in this area face a gap between these concepts and the question of how they can best be integrated in a systematic and unified way. One of the main reasons for this gap is the fact that the core concepts in perception, action, and communication are typically studied by different communities: computer vision, robots, and natural language-processing communities, among others, without much interchange between them. As machine learning lies at the core of these communities, it can act as a unifying factor in bringing the communities closer together. Unifying these communities is important for understanding how state-of-the-art approaches from different disciplines can be combined and applied to build generally intelligent interactive systems.

This workshop brought together researchers from a variety of subfields of AI such as robotics, machine learning, natural language processing, and human-robot interaction. The papers presented at the workshop addressed the following machine-learning frameworks: supervised learning, unsupervised learning, reinforcement learning, active learning, transfer learning, and learning from demonstration. While the major theme of papers presented at the workshop was interactive robots, interactive games and dialogue systems were other very well received themes. The workshop also had poster presentations, which were useful to encourage discussion among participants. In addition, the workshop featured four invited talks given by renowned researchers. Andrea Thomaz (Georgia Institute of Technology, USA) gave a talk on socially interactive robot learning. Joelle Pineau (McGill University, Montréal, Canada) spoke on the lessons learned from the SmartWheeler project — learning socially adaptive navigation strategies. Sonia Chernova (Worchester Polytechnic Institute, USA) presented a talk about enabling robots to learn from everyday people. Finally, Matthias Scheutz (Tufts University, USA) presented work on natural language-guided one-shot learning.

The workshop concluded with a discussion session that included panelists Andrea Thomaz, Sonia Chernova, Matthias Scheutz, and George Konidaris. The panelists engaged in conversations with the audience regarding the progress made on a number of topics, namely, in bridging the gap between perception,

action, or communication; the limited success stories of interactive learning systems; the need for publicly available data and tools for multimodal interactive systems; the crucial need for understanding the principles of human-machine interaction to come up with a breakthrough in intelligent interactive systems; and, finally, the need to hold focused events of this kind to continue discussing issues that affect the research, development, and deployment of learning interactive systems.

Heriberto Cuayáhuitl, Lutz Frommberger, Nina Dethlefs, and Martijn van Otterlo wrote this report and served as cochairs of this workshop. The papers of the workshop were published by AAAI Press, Technical Report WS-14-07.

Modern Artificial Intelligence for Health Analytics

The proliferation of healthrelated data presents unprecedented opportunities to improve patient care, but medical experts are currently overwhelmed by information. Therefore, novel artificial intelligence approaches are needed to process, organize, and transform these data into actionable knowledge. Building on existing AI techniques, developing new ones, and evaluating such methods will require interdisciplinary, collaborative research between health sciences professionals and computer scientists.

Another challenge is that health information is extremely heterogeneous, including, for example, results from clinical trials, published biomedical literature, patient health records, genomic data, wearable health monitor outputs, online reviews of physicians, and medical images. Transforming these data in a meaningful way requires innovations in methods spanning the whole of AI, from natural language processing to computer vision. Working with health data poses many other challenges: the data tend to be sparse, are often unsemistructured or semistructured, and are noisy.

This workshop sought to bring together artificial intelligence researchers working in this interdisciplinary space to discuss progress, challenges, and opportunities. The workshop featured four invited talks from distinguished speakers. Eric Horvitz (Microsoft Research) outlined promising directions for leveraging machine intelligence to improve health care in a datarich era. Carla Brodley (Northeastern University) presented recent work on using hierarchical structured models for brain lesion detection in medicationresistant epilepsy patients. Rich Caruana (Microsoft Research) presented work on the importance of transforming machine-learning outputs into wellcalibrated probabilities for health-care data. Peter Szolovits (MIT) spoke about AI for the realization of a learning health-care system and presented recent work on learning to predict adverse events using critical care data.

Additionally, we had several presentations based on submitted work. These works covered a diverse set of health applications and methodological approaches including sequential batch reinforcement learning for personal medical decision making; models for fMRI classification and biomarker discovery; prediction of blood glucose levels for diabetes management; crowdsourcing a clinical trial repository for evidencebased medicine; and optimization strategies for kidney exchanges.

The workshop concluded with a panel discussion, led by Christopher Yang (Drexel University) and featuring the invited speakers. The panel discussed concerns pertaining to the emerging community of researchers at the intersection of AI and health. This discussion led to the topic of nascent publication venues such as the Meaningful Use of Complex Medical Data conference and various highquality workshops that have been appearing at major conferences such as AAAI. There was consensus that these venues would persist but undoubtedly grow in size as the community itself continues to grow, and that encouraging more crosstalk between computer scientists and health sciences researchers is important.

Byron C. Wallace and Jenna Wiens wrote this report. This workshop was organized by Byron Wallace, Jenna Wiens, David Kale, and Finale DoshiVelez. The workshop proceedings were published as AAAI Press Technical Report WS1408.

Multiagent Interaction Without Prior Coordination

Interaction between agents is the defining attribute of multiagent systems, encompassing problems such as planning in a decentralized setting, learning 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, the majority of these methods depend on some form of prior coordination. For example, learning algorithms may assume that all agents share a common learning method or prior beliefs, distributed optimization methods may assume specific structural constraints regarding the partition of state space or cost/rewards, and symbolic methods often make strong assumptions regarding norms and protocols. However, in realistic problems these assumptions are easily violated; hence 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 eight accepted papers, with topics as diverse as agent coor-

dination through program equilibrium, ad hoc teamwork in robot soccer, communication-restricted exploration for robot teams, and multiagent learning in highly dynamic environments. The presented research demonstrated that MIPC problems exist in various flavors and that there are a variety of approaches to tackle such problems.

We were privileged to have invited talks by three distinguished researchers with background in MIPC: Subramanian Ramamoorthy from the University of Edinburgh gave a talk on a categorization-based approach to ad hoc human-machine interaction, Peter Stone from the University of Texas at Austin gave a talk on collaboration without precoordination of ad hoc autonomous agent teams, and Manuela Veloso from Carnegie Mellon University gave a talk on synergy graphs to train and form ad hoc teams.

The final highlight of the workshop was an expert panel consisting of Noa Agmon from Bar-Ilan University, Prashant Doshi from the University of Georgia, and Michel Littman from Brown University. The panelists discussed their opinions on important issues in MIPC research and new developments in this emerging area. Following opening statements from each panelist, the workshop participants were invited to ask questions. This resulted in a lively and interesting conversation.

This report was written by Stefano V. Albrecht. The workshop was chaired by Stefano Albrecht, Samuel Barrett, Jacob Crandall, and Somchaya Liemhetcharat. The chairs would like to thank the workshop participants, the invited speakers, the expert panelists, the program committee, and the AAAI staff for making the workshop a success. The papers of the workshop were published as AAAI Press Technical Report WS-14-11.

Multidisciplinary Workshop on Advances in Preference Handling

Preferences are inherently a multidisciplinary topic and are fundamental for the analysis of human choice behavior and in designing intelligent agents that act on our behalf. They are becoming of increasing importance for computational fields such as artificial intelligence, databases, and human-computer interaction. The goal of this workshop was to address diverse computational aspects of preference handling and to provide cross-fertilization between different fields.

This workshop brought together researchers from a variety of subfields of artificial intelligence and database systems, such as decision support, multiagent systems, social choice, recommender systems and elicitation, learning, modeling, representation, aggregation, and management of preferences. The workshop presentations covered four general topics. Talks on preference representation included papers on CP-networks, utility hypergraphs, and other preference formalisms used in AI and database systems. The preference reasoning session included papers about reasoning under the Pareto principle, the likelihood of structures in preference profiles, and aggregating opinions to design energy-efficient buildings. Another general topic was on group-based decision making and, specifically, analysis of different voting schemes, such as approval-based committee voting, or multiwinner approval voting. Finally, the concluding session served to remind the participants that preferences can also be used in playing chess, eating, or recommendations. The workshop included an invited talk by Craig Boutilier (University of Toronto) on scaling optimization methods for data-driven marketing. The broadened scope of the preference workshop should enable closer future collaboration of researchers from distinct fields working on solving similar problems.

Markus Endres, Darius Braziunas, and Paul Weng wrote this report. The workshop was organized by Darius Braziunas, Markus Endres, K. Brent Venable, Paul Weng, and Lirong Xia. The papers of the workshop were published as AAAI Press Technical Report WS-14-10.

Semantic Cities — Beyond Open Data to Models, Standards, and Reasoning

In a semantic city, available resources are harnessed safely, sustainably, and efficiently to achieve positive, measurable economic and societal outcomes. Enabling city information as a utility, through a robust, expressive, dynamic, scalable, and, critically, sustainable technology and socially synergistic ecosystem, could drive significant benefits and opportunities. Data, and then derived information and further derived knowledge, from people, systems, and things, is the single most scalable resource available to city stakeholders to reach the objective of semantic cities.

Two major trends are supporting semantic cities — open data and semantic web. "Open data is the idea that data should be accessible from everyone to use and republish as they wish, without restrictions from copyright, patents, or other mechanisms of control." A number of cities and governments have made their data publicly available, prominent being London (UK), Chicago (USA), Washington, D.C. (USA), and Dublin (Ireland).

Semantic web as the technology to interconnect heterogeneous data has matured and it is being increasingly used in the form of linked open data and formal ontologies. Thus, a playfield for more AI research-driven technologies for cities has emerged, for example, scalable, efficient, robust, optimal AI techniques.

We had a rich program highlighting how to build

semantic cities and provide societal benefits. It consisted of an invited talk, seven regular papers, four short papers, and demonstrations (presented as posters), and a panel attracting speakers and attendees from academe and industry. We discussed ontologies, their design, expressivity, and potential impact for semantic cities — AI topics such as machine learning, data mining, and optimization; use cases related to semantic-augmented health care; public services such as 311; and citizen-driven applications, social care, ecology, e-governance.

Questions related to the level of city data representation and, interestingly, ontologies, for example, their expressivity, extension, and representation, were discussed, and the participants finally concluded that the trade-off between expressivity and tractability needs to be addressed case by case. Any generalization is then very difficult to be reached. The panel pointed out important questions such as the replicability of research works among different cities and open access to data sets. Mark Fox pointed out that cities can be smarter only if we can assess and compare them on various major metrics.

The workshop participants discussed how interest in semantics for cities and public bodies is growing while pointing out the lack of best practices to reach the level of semantic cities. In this respect the participants agreed that the workshop was useful and successful. Specially, they really liked the application-driven presentation of various contributions of the workshop even if some participants asked for in-depth technical details. Some participants suggested a hackathon around some available data sets from major cities. They also clearly agreed on the needs and benefits of applying and revisiting AI techniques for solving real-world city problems.

The latter was appreciated as they saw the potential that new research challenges arise from these real-world domains. The multidisciplinary dimension of the contribution was also appreciated — the fusion of machine learning, optimization, data mining, and knowledge representation for tackling problems city managers are facing every day.

Mark Fox, Freddy Lecue, Sheila McIlraith, Biplav Srivastava, and Rosario Uceda-Sosa wrote this report and served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-13-01.

Sequential Decision Making with Big Data

In the 21st century, we live in a world where data is abundant. We would like to use this data to make better decisions in many areas of life, particularly those with significant social impact, such as industry, health care, business, and government. This opportunity has encouraged many machine-learning and data mining researchers to develop methods to benefit from big data. However, the methods developed so far have focused almost exclusively on the task of prediction. As a result, the question of how big data can help in the solution of complex decision-making problems, especially sequential ones, has remained largely untouched. The goal of the workshop was to discuss how the current availability of massive amounts of data can help solve large-scale decision-making problems, especially those involving a sequence of choices whose effects accumulate over time.

The workshop brought together around 40 researchers from different subfields of artificial intelligence, in particular from the reinforcement learning community. Many interesting topics were discussed, including new reinforcement learning techniques, challenging sequential decision-making benchmarks, and important real-world applications. There was also some emphasis on the multiarmed bandit problem, a simplified yet important version of the sequential decision-making scenario.

The workshop had four invited talks. In the first talk Richard Sutton (University of Alberta) described a new version of the temporal-difference learning algorithm whose conceptual description corresponds exactly to its practical implementation. In the second invited talk Emma Brunskill (Carnegie Mellon University) discussed how several challenges that arise in education can be modeled and solved as decision-making problems. Lihong Li (Microsoft Research) described in the third invited talk the successes, lessons, and challenges of using multiarmed bandits to solve decision-making problems on the Internet. In the fourth and final invited talk Michael Bowling (University of Alberta) presented the Arcade Learning Environment, a video game emulator that serves as a test bed for the development and evaluation of general agents.

In addition to the invited talks, there were six contributed talks presenting the papers accepted in the workshop. The topics discussed include exploration strategies for robotics, feature extraction and representation learning in reinforcement learning, and novel applications of multiarmed bandits. The workshop ended with an hour-long roundtable in which many members of the audience engaged in fruitful discussions related to the topics considered earlier.

Overall the workshop was very successful in bringing the community together and inaugurating the discussion on sequential decision making with big data. This initial effort has the potential of unfolding into other events on the same theme.

André M. S. Barreto and Amir-massoud Farahmand wrote this report. Joelle Pineau and Doina Precup, from McGill University, and Mohammad Ghavamzadeh, from Adobe Research and INRIA Lille, served as cochairs of the workshop. The papers of the workshop were published as AAAI Technical Report WS-14-12.

Statistical Relational AI

Much has been achieved in the field of AI, "the science and engineering of making intelligent machines" as John McCarthy defined it, yet much remains to be done if we are to reach the goals we all imagine. One of the key challenges moving ahead is closing the gap between logical and probabilistic AI. Logical AI has mainly focused on complex representations, and probabilistic AI on uncertainty. Intelligent agents, however, must be able to handle both the complexity and the uncertainty of the real world. Combining logic and probability in a unified representation and building general-purpose reasoning tools for such representations has been the dream of AI, dating back to the work on probabilistic logic by Nils Nilsson published in the *AI Journal* in 1986. Despite the fact that the two fields share many key features and often solve similar problems, research in them has progressed independently with little interaction, and sharing results has been rather cumbersome. The Statistical Relational AI workshop was designed for attempts at synthesis, toward a big-picture view on AI, and ultimately to explore what might be called statistical relational AI: "the science and engineering of making intelligent machines that act in uncertain environments composed of objects and relations among the objects".

The last decade has seen real progress toward closing the gap, with an explosion of successes in combining probability and (subsets of) first-order logic, programming languages, and relational databases in several subfields of AI such as reasoning, learning, knowledge representation, planning, databases, natural language processing, robotics, vision, and others. Today, we can learn statistical relational models automatically from millions of interrelated objects. We can generate plans and learn to act optimally in uncertain environments involving objects and relations among them. We can even perform lifted probabilistic inference, avoiding explicit state enumeration by manipulating first-order representations directly. So far, however, the researchers combining logic and probability in each of these subfields have been working mostly independently.

This workshop convened researchers that drive forward different AI subfields by using statistical relational techniques. In the first part, the invited talks by Luc De Raedt (Katholieke Universiteit Leuven), Vibhav Gogate (University of Texas, Dallas), and Henry Kautz (University of Rochester) provided a synthesis of probabilistic programming and lifted inference and a statistical relational AI perspective on recognizing complex events. The second part of the workshop was a lively poster session where 14 full technical papers and 13 short position papers and abstracts were presented. The participants were encouraged to discuss the commonalities and differences among the various AI tasks addressed. The selected papers covered a wide range of topics in statistical relational AI, such as parameter and structure learning, lifted inference, exchangeability, probabilistic programming, planning, information extraction, constraint optimization, natural language processing, among others, and clearly show the promise of statistical relational techniques for AI. One of the main themes of the workshop was scalability in the face of complex probabilistic interactions, and lifted probabilistic inference in particular. The group reached a general consensus that statistical relational AI is an exciting emerging area that requires more investigation.

Guy Van den Broeck, Kristian Kersting, Sriraam Natarajan, and David Poole served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-14-13.

The World Wide Web and Public Health Intelligence

The aim of the AAAI workshop on the World Wide Web and Public Health Intelligence was to investigate the theory and practice of computational models of web-based public health intelligence to highlight the latest achievements in epidemiological surveillance based on monitoring online communications and interactions on the Web.

Public health Intelligence aims to promote health, prevent/reduce disease, and provide appropriate timely response/alert to critical public health events (that is, disease outbreaks, bioterrorism, and others) through intelligent knowledge exchange and real-time surveillance. It also provides platforms that collect, integrate multiple data sources, and apply analytics to generate insights that improve decision making at individual and societal levels. Given the ever increasing role of the World Wide Web as a source of information in many domains including health care, accessing, managing, and analyzing its content has brought new opportunities and challenges. Rapid increase in access to the Internet and widespread use of online social media along with the massive amount of online streaming data generated by hundreds of millions of devices, users, service providers, government agencies, and other stakeholders have made the effective use of these data an emerging priority for public health intervention. The goal of this workshop was to provide a forum for knowledge exchange among a multidisciplinary and international crowd of researchers, students, professionals, public health agencies, and others interested in the theory and practice of computational models of web-based public health intelligence.

Presentations highlighted the latest achievements in epidemiological surveillance based on monitoring online communications and interactions on the web. The workshop brought together researchers from a variety of subfields of artificial intelligence such as machine learning, knowledge representation, natural

language processing, online social media analytics, ontology engineering, and decision-support systems in addition to researchers in fields of epidemiology, public health informatics, disease surveillance, and emergency management. One major theme of papers presented at the workshop was analyzing online behaviors, social media contents, and online health discussions for determining patterns to enhance forecasting, decision making, and classification. The workshop also included an invited talk on this theme given by Eric Horvitz (Microsoft Research), which focused on harnessing anonymized behavioral data drawn from online services as large-scale sensor networks for health and well-being. Horvitz discussed opportunities for garnering insights and performing inferences and predictions about health through analyses of signals expressed in social media, web searching and browsing, and interpersonal communications.

To promote open debate and exchange of opinions among participants, the workshop closed with a panel discussion moderated by David L. Buckeridge (McGill University) and including Eric Horvitz (Microsoft Research), Denis A. Roy (Institut National de la Santé Publique du Québec), Reiner Banken (Institut National d'Excellence en Santé et en Services Sociaux). The theme of the panel was the requirements, as well as opportunities and barriers, for translating advances in public health intelligence into practice.

Arash Shaban-Nejad, David L. Buckeridge, and John S. Brownstein served as cochairs of this workshop. The papers of the symposium were published as AAAI Press Technical Report WS-14-14.

Notes

1. sites.google.com/site/cognitivecomputingahi

Stefano V. Albrecht is a Ph.D. candidate in the School of Informatics at the University of Edinburgh.

André M. S. Barreto is an assistant researcher in the Department of Applied and Computational Mathematics at the Brazilian National Laboratory for Scientific Computing.

Darius Braziunas is a senior research scientist in the Big Data group at Kobo Inc.

David L. Buckeridge, MD, PH.D., is an associate professor in the Department of Epidemiology and Biostatistics at McGill University.

Heriberto Cuayáhuitl is a senior research fellow in the School of Mathematical and Computer Sciences at Heriot-Watt University in Edinburgh, Scotland.

Nina Dethlefs is a research fellow in the School of Mathematical and Computer Sciences at Heriot-Watt University in Edinburgh, Scotland.

Markus Endres is a researcher and lecturer at the University of Augsburg, Germany.

Amir-massoud Farahmand is an NSERC postdoctoral fellow currently with Carnegie Mellon University.

Mark Fox is a professor at the enterprise integration laboratory at the University of Toronto, Canada.

Lutz Frommberger is a senior researcher in the Cognitive Systems Group at the University of Bremen.

Sam Ganzfried is a Ph.D. student in the computer science department at Carnegie Mellon University.

Sébastien Guillet is a postdoctoral research fellow in computer science at LIARA (UQAC, Chicoutimi, Canada).

Yolanda Gil is director of knowledge technologies and associate division director at the Information Sciences Institute of the University of Southern California, and research professor in the Computer Science Department.

Lawrence E. Hunter is director of the Computational Bioscience Program and the Center for Computational Pharmacology at the University of Colorado School of Medicine, and a professor of pharmacology, biology, biometrics, and computer science.

Arnav Jhala is an assistant professor in the Department of Computational Media at the University of California Santa Cruz.

Kristian Kersting is an associate professor with the Computer Science Department at the Technical University of Dortmund, Germany.

George Konidaris is a postdoc at the MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, Massachusetts.

Freddy Lecue is a research scientist and lead investigator in large-scale reasoning systems at IBM Research – Smarter Cities Technology Centre, Dublin, Ireland.

Sheila McIlraith is a professor at the department of computer science at the University of Toronto, Canada.

Sriraam Natarajan is an assistant professor at the School of Informatics and Computing, Indiana University, USA.

Zeinab Noorian is a postdoctoral research fellow in the Department of Computer Science at the University of Saskatchewan, Canada.

David Poole is a professor of computer science at the University of British Columbia, Canada.

Rémi Ronfard is a senior researcher at INRIA and the University of Grenoble, France.

Alessandro Saffiotti is a full professor of computer science at Orebro University, Sweden.

Arash Shaban-Nejad, Ph.D., is a postdoctoral fellow in the McGill Clinical and Health Informatics group at McGill University.

Biplav Srivastava is a senior research scientist at IBM Research.

Gerald Tesauro is a research staff member in the Cognitive Computing Department at IBM Research, Yorktown Heights, New York.

Rosario Uceda-Sosa is currently one of the leading researchers in semantic information models for Smarter Cities in IBM Research

Guy Van den Broeck is a postdoctoral researcher at the Computer Science Department, Katholieke Universiteit Leuven, Belgium.

Martijn van Otterlo is a senior researcher in the Cognitive Artificial Intelligence Department of the faculty of Social Sciences at Radboud University Nijmegen.

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

Save the Date for AAAI-16

Mark your calendars now for the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16) and the Twenty-Eighth Conference on Innovative Applications of Artificial Intelligence (IAAI-16), which will be held in sunny Phoenix, Arizona, February 12-17, 2016! Please note the change in day pattern from the traditional AAAI model. Workshops and tutorials will be held Friday and Saturday, February 12-13, with the main conference starting on Sunday, February 14. The technical conference will run 3.5 days, as it did in 2015. Phoenix is the gateway to the Grand Canyon, and its history is a testament to the spirit of puebloans, ranchers, miners, and visionaries. Projected against this rich backdrop is a panorama of urban sophistication, with a host of museums (be sure to visit the Pueblo Grande Museum and Archaeological Park and the Heard Museum), sports stadiums, restaurants, and shopping. Nearby Tempe is the site of Arizona State University, home of a leading AI research community. The 2016 program cochairs are Dale Schurmanns and Michael Wellman. The Call for Papers will be available soon (www.aaai.org/aaai16). Please join us in 2016 in America's sunniest metropolis!

Paul Weng is an associate professor in computer science at the Pierre and Marie Curie University, Paris, France.

Jenna Wiens is an assistant professor at the University of Michigan.

Jie Zhang is an assistant professor at the School of Computer Engineering, Nanyang Technological University, Singapore.