Reports of the 2014 AAAI Spring Symposium Series

Manish Jain, Albert Xin Jiang, Takashi Kido, Keiki Takadama, Eric G. Mercer, Neha Rungta, Mark Waser, Alan Wagner, Jennifer Burke, Don Sofge, William Lawless, Mohan Sridharan, Nick Hawes, Tim Hwang

■ *The Association for the Advancement of Artificial Intelligence was pleased to present the AAAI 2014 Spring Symposium Series, held Monday through Wednesday, March 24–26, 2014. The titles of the eight symposia were Applied Computational Game Theory, Big Data Becomes Personal: Knowledge into Meaning, Formal Verification and Modeling in Human-Machine Systems, Implementing Selves with Safe Motivational Systems and Self-Improvement, The Intersection of Robust Intelligence and Trust in Autonomous Systems, Knowledge Representation and Reasoning in Robotics, Qualitative Representations for Robots, and Social Hacking and Cognitive Security on the Internet and New Media). This report contains summaries of the symposia, written, in most cases, by the cochairs of the symposium.*

Applied Computational Game Theory

Game theory's popularity continues to increase in a variety of disciplines such as economics, biology, political science, computer science, electrical engineering, business, law, public policy, and many others. The focus of this symposium was to bring together the community working on applied computational game theory motivated by any of these domains. This symposium, while not limited to the ideas discussed there, built on the AAAI Spring Symposium 2012 on Game Theory for Security, Sustainability, and Health. With the development of new computational approaches to game theory over the past two decades, very large-scale real-world problems can be cast in game-theoretic contexts and solved efficiently, thus providing us computational tools to address real-world problem sizes. For instance, in the arena of security, there now seems to be an exponential increase in interest

 ISSN 0738-4602

due to the emergence of computational game theory. Software assistants have been developed for randomized patrol planning for the Los Angeles International Airport police, the Federal Air Marshal Service, United States Coast Guard, and the Los Angeles Sheriff's Department.

This symposium brought together researchers from a variety of subfields of AI, including behavioral modeling, computational game theory, multiagent robotics, public policy, operations research, planning, and machine learning. The presentations varied among all these domains, while the approaches spanned algorithmic and computational game theory, decision theory, and machine learning, as well as Markov decision processes (MDPs). We had 7 invited speakers and 10 student presentations at the symposium. Student presentations ranged from most recent results on patrolling metro transit systems, game-theoretic scheduling of defender resources in conservation domains like forests and fisheries, computing deceptive strategies as well as developing robust algorithms that handle all forms of uncertainty in security domains. Other student presentations covered topics like identifying when to kick out cheating gamblers, using user incentives to reduce energy consumption, and using MDPs to coordinate resources. We also had a presentation covering theoretical properties of Stackelberg equilibria.

The invited talk by Kevin Leyton-Brown from the University of British Columbia discussed human behavior — both designing human subject experiments as well as predictive behavioral models. Another invited talk by Tuomas Sandholm from Carnegie Mellon University described abstracting very large games and showed most recent results on computing strategies for large extensive-form games like poker. We had invited talks that discussed most recent research and results in adversarial robotics (Noa Agmon, Bar-Ilan University), disaster management (Jun Zhuang, University of Buffalo), cybersecurity and strategic placement of honeypots (Christopher Kiekintveld, University of Texas at El Paso), and adversarial machine learning and spam detection (Eugene Vorobeychik, Vanderbilt University). Another invited talk by Sam Barrett (University of Texas at Austin) described the latest research in ad hoc teamwork.

Overall, the symposium participants discussed the different avenues of game-theoretic research, and the symposium was useful in bringing together researchers from many different subfields with many different perspectives. The participants engaged in enthusiastic discussions on the topics and expressed excitement over attending future symposia with the same (and even broader) focus as this one.

Manish Jain, Alberg Xin Jiang, Bo An, and Samarth Swarup served as cochairs of this symposium. No technical report for this symposium was published.

Big Data Becomes Personal: Knowledge into Meaning

The goal of the Big Data Becomes Personal: Knowledge into Meaning symposium was to explore how big data can be made personally usable and meaningful using AI technologies.

This symposium builds on the 2013 AAAI Spring Symposium on Data-Driven Wellness: From Self-Tracking to Behavior Change. Because of the participants' strong desire to continue the event we organized it again, this time incorporating the new concept: big data becomes personal. We think that one of the most significant shifts in our contemporary world is the trend toward obtaining and analyzing big data in nearly every venue of life. For better health, wellness, and well being, it is very significant to acquire personal meaningful information from big data. However, the following outstanding challenges should be tackled to make big data become personal: (1) How do we quantify our health, wellness, and well-being for generating useful big data that will become meaningful knowledge? (2) How do we turn the large volumes of impersonal quantitative data into qualitative information that can affect the quality of life of the individual? (3) How do the quantitative data and qualitative information contribute to improving our health, wellness, and well-being? This symposium explored the methods and methodologies for the above three questions.

Two special invited speakers gave us new perspectives. First, Rosalind W. Picard (MIT Media Lab) discussed her affective computing project, describing how sensors can show surprising and important meaning. Second, John Perry (Stanford University) spoke about personal identity, which closely relates to realistic issues in future personal big data societies. For example, mind uploading (sometimes called mind copying or mind transfer), which copies our mental content (including long-term memory and self) from a particular brain substrate to another computational device, presents us with many new philosophical questions about our personal identity.

The symposium also included 23 paper presentations and 5 poster presentations, organized into 11 topics: (1) detection and prediction from big data, (2) cognitive health, (3) care worker support system, (4) big data analysis, (5) vital and brain data monitoring and analysis, (6) mobile and self-tracking health data, (7) sensing data for health and wellness, (8) personal identity, (9) detection from sensor data, (10) emotion and interface, and (11) personal big data and repository.

Some highlights of the presentations included Mihoko Otake (Chiba University), who introduced her field work on dementia projects. Her analyses with cognitive information theory showed that interactive group discussions could improve the cognitive health. Ida Sim (University of California, San Fran-

cisco) introduced her Open mHealth project and discussed the open architecture for integrating mobile data for health care. Akane Sano (MIT Media Lab) introduced her affective computing project research on human sleep and discussed the issues of understanding ambulatory and wearable data for health and wellness. J. T. Turner (University of Maryland) introduced analytical research on the use of a variety of representations and machine-learning algorithms applied to the task of seizure detection in high-resolution, multichannel EEG data. Koichi Hasida (National Institute of Advanced Industrial Science and Technology, Japan) gave a talk on a new platform for a personal life repository that allows individual users to totally control their own data and drastically reduces the cost and risk to service providers in storing personal data. Yukio Ohsawa (Tokyo University) presented the idea of an innovators marketplace where the content of each data set may be hidden due to constraints that are strict in health care but the digest of the data set can and should be disclosed for expressing the latent value of the data.

One of the important outcomes of our symposium was to present a new perspective from the viewpoint of personal big data. Since personal data, such as health-care data, is related to ourselves, it motivates us. For example, we might start running if we find ourselves gaining weight, even if we do not usually run for our health. This indicates that the personal data has far greater power to implement behavior change in comparison to nonpersonal data. Another aspect of personal big data is security. Since the personal data is private, such data should be secured. However, it is hard to protect all data as the volume or type of data increases. To address this issue, a (distributed) personal data store was found to be a core approach for personal big data. These are but two unique aspects of personal big data. We expect to continue discussing these important interdisciplinary challenges for guiding future advances in the AI community.

Takashi Kido and Keiki Takadama served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-14-01.

Formal Verification and Modeling in Human-Machine Systems

Recent papers in formal verification, cognitive modeling, task analysis, and human-machine systems discuss modeling challenges and the application of basic formal verification in human-machine interaction. This symposium brought together experts from many communities to explore key research areas, common solutions, near-term research problems, and advantages in combining the best of the different communities.

An important theme of the symposium was the challenges associated with modeling human-machine systems. A panel session, titled Modeling Hurts, identified some of the challenges associated with creating and using models, including the time to create the models and the inherent limitations of the models, but all panelists agreed that good but flawed models were necessary for understanding and improving human-machine interaction. Several papers were directly relevant to this theme and discussed particular models for particular problems. The significant lesson from these papers was partly that fundamental work still needs to be done to create useful models of humans and machines interacting, and partly that there is no way to efficiently integrate different modeling approaches to create something that is better than any individual model alone can achieve.

A second important theme was the somewhat surprising observation that when talking about using verification and modeling in human-machine systems, a meta-human-machine interaction emerges, namely that of providing support for the human responsible for creating and verifying the models. Issues related to supporting the modeler include supporting a range of abstraction levels and providing automated support for what was called rapid modeling, an intentional reference to the field of rapid prototyping and its potential for improving the design process.

A third important theme was the practical limitations imposed by the complexity of using model checking and formal methods for verifying human-machine systems. Managing this complexity requires reliable models with known limits and manageable complexity so that robust results can be generated using established methods. An important area of future work is how to choose the right level of abstraction for a given problem, with the goals of the modeling and evaluation effort to produce reasonable bounds on system performance including the ability to know when a system is operating outside a safe envelope without resorting to the unacceptable extremes of treating the human as an infallible oracle, on one hand, or the human as a disturbance to be rejected, on the other.

Neha Rungt, Eric Mercer, Ellen Bass, and Michael Goodrich served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-14-02

Implementing Selves with Safe Motivational Systems and Self-Improvement

While artificial intelligence (AI) and artificial general intelligence (AGI) most often focus on tools for collecting knowledge and solving problems or achieving goals rather than self-reflecting entities, this implementation-oriented symposium focused

on guided self-creation and improvement — particularly as a method of achieving safe and ethical human-level intelligence in machines through iterative improvement (seed AI).

The first day focused on creating self-improving selves and was led off by invited speaker Daniel Silver's (Acadia University) talk on lifelong machine learning and reasoning. Themes included lifelong learning, incorporating the self into active logic, failures common to autistic humans and learning systems, and languages and methods for serving up minds. Invited speaker Pierre-Yves Oudeyer (Flowers Laboratory, Inria and Ensta ParisTech) continued the topic on the third day with his talk on developmental robotics, lifelong learning, and the morphogenesis of developmental structures.

The second day focused on safe motivational systems with speakers presenting ideas ranging from tying AGIs to human emotions and values (either directly or through a consensus harvested from the web) to designing a humanlike motivational structure to methods for constraining selves to safe behavior. There were also several robust discussions around a number of the currently popular memes, both pro and con. The first invited speaker, Daniel Polani (University of Hertfordshire) talked about Empowerment: A Universal Utility for Intrinsic Motivation or What to Do When You Do Not Know What to Do? while Steve Omohundro (Self-Aware Systems) discussed ethics and understanding and managing the unintended consequences of self-improvement.

The general opinion of the participants, with notable exceptions, was that the field is not yet to the point where it is feasible actually to attempt to implement a complete safe self-improving artificial entity. All agreed, however, on the critical importance of achieving that goal in a timely fashion and that they would like to attend future symposia with the same focus as this one.

Mark Waser served as chair of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-14-03.

The Intersection of Robust Intelligence and Trust in Autonomous Systems

The AAAI symposium on the intersection of robust intelligence and trust developed from the interest and excitement generated from last year's symposium on trust in autonomous systems. The symposium was meant to examine the intersection of robust intelligence and trust across multiple environments among autonomous hybrid systems (where hybrids are arbitrary combinations of humans, machines, and robots). The symposium investigated methods for structuring teams or networks of computers that would increase robust intelligence and engender trust among a system of agents (people or machines).

One goal of the symposium was to better understand robust intelligence and the autonomy it produces among humans interacting with other humans and human groups (for example, teams, firms, systems; also, networks among these social objects). Ultimately, we believe that the information divined from a better understanding of robust intelligence will not only be useful to artificial intelligence researchers but will also allow us to predict the interactive outcomes of hybrid human-machine groups.

Systems that learn, adapt, and apply experience to problems may be better suited to respond to novel environmental challenges. One could argue that such systems are robust to the prospect of a dynamic and occasionally unpredictable world. We expect systems that exhibit robustness would afford a greater degree of trust to those people that interact with the system. Robustness through learning, adaptation, and knowledge is determined by predicting and modeling the interactions of autonomous hybrid systems. How can we use this data to develop models indicative of normal/abnormal operations in a given context? We hypothesize such models improve system intelligence by allowing autonomous entities to adapt continually within normalcy bounds, leading to greater reliability and trust.

The focus of the symposium was on how robust intelligence affects trust in the system and how trust in the system affects robustness. We explored approaches to robust intelligence and trust employing a variety of different kinds of vehicles, environments, and purposes. Speakers presented applications that ranged from methods for optimizing satellite imagery to interacting with one's Facebook friends. The symposium explored theory, mathematics, computational models, and field applications at the intersection of robust intelligence and trust. Side discussions debated the meaning of robust intelligence and how the two topics relate to one another. Satyandra Gupta, a program director with the National Science Foundation and professor of mechanical engineering, noted that robust intelligence researchers tend to utilize machine learning as an approach for increasing a system's robustness and adaptability. However, leaders in the military tend to avoid machines that learn because of a lack of trust. Hence, one could argue that the methods used to make a system robust could simultaneously decrease trust in the system or, alternatively, that learning by single agents penalizes team users. Suzanne Barber, an endowed professor of electrical engineering with University of Texas, presented the challenges and implications of robustly verifying one's identity and the impact that this can have on trusted interactions in cyberspace. The inchoate nature of the topics led to many lively discussions.

The symposium was well attended. Approximately five to six nonspeaking additional participants also regularly attended the sessions. Application areas related to trust and robust intelligence were a reoccurring theme for the papers presented at the symposium. A presentation by Sean Augenstein (Skybox Imaging) considered the trust in personalized satellite imagery. Hadas Kress-Gazit (Cornell University) discussed verification of high-level robot behavior and its impact on robustness. Future plans for this research area include follow-on meetings and possibly a more thorough article on the topic developed for journal publication.

Alan Wagner, Jennifer Burke, Don Sofge, and William Lawless served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report SS-14-04.

Knowledge Representation and Reasoning in Robotics

Robots deployed in real-world application domains face the formidable challenge of representing, learning from, and reasoning with incomplete domain knowledge acquired from sensors, humans, and other sources. Although many algorithms and architectures have been developed for qualitatively or quantitatively representing and reasoning with knowledge, the research community is fragmented, with separate vocabularies that are increasingly (and ironically) making it difficult for these researchers to communicate with each other. For instance, the commonsense reasoning capabilities of declarative languages are not well understood or fully used by the robotics community, while the logic programming community does not fully appreciate the ability of probabilistic algorithms to elegantly model the uncertainty in sensing and actuation on robots. The objective of this symposium was to bring these disparate groups together to promote a deeper understanding of recent breakthroughs and open problems in the logic programming and probabilistic reasoning communities, thus encouraging collaborative efforts toward building architectures and algorithms that support representation and reasoning with qualitative and quantitative descriptions of knowledge and uncertainty in robotics.

The symposium program included paper presentations on topics of interest such as knowledge representation, commonsense reasoning, reasoning about uncertainty and with incomplete knowledge, interactive and cooperative decision making, symbol grounding, and combining symbolic and probabilistic representations. Some papers described novel theoretical contributions, while other papers grounded and illustrated the algorithms in research areas such as robot vision, robot planning, and human-robot (and multirobot) collaboration.

The symposium had five invited talks that described long-term research in representation and reasoning for robots and agents. The sessions with the invited speakers, and the morning session on the last day, were shared with a parallel symposium on Qualitative Representations for Robots. On the first day, Jeffrey Siskind (Purdue University) gave a talk on the compositional structure of perception, language, action, and thought, while Daniele Nardi (Sapienza University of Rome) gave a talk on semantic mapping. On the second day, Anthony Cohn (University of Leeds) described efforts to learn qualitative spatiotemporal activity models, while Michael Gelfond (Texas Tech University) described P-log, a knowledge representation language capable of representing logical and probabilistic reasoning. On the final day, Mary-Anne Williams (University of Technology, Sydney) gave a talk on representation and the role of attention for social robots.

The paper presentations and invited talks promoted discussions and exchange of ideas among the participants and helped establish connections that we hope will lead to long-term collaborations. The participants stated that the symposium was an excellent learning experience and expressed a strong interest in attending future symposia on knowledge representation and reasoning in robotics.

Mohan Sridharan, Fangkai Yang, Subramanian Ramamoorthy, Volkan Patoglu, and Esra Erdem served as cochairs of this symposium. The papers presented at the symposium were published as AAAI Press Technical Report SS-14-05.

Qualitative Representations for Robots

As the field of robotics matures, the construction of ever more intelligent robots becomes possible. For many of the challenging tasks we want robots to perform it is crucial that a robot can be provided with knowledge: knowledge of its capabilities, of its environment, and of how the former interacts with the latter. The fields of AI and robotics have many approaches to representation and reasoning. This symposium focused on one approach that has been growing in popularity across these communities in recent years: qualitative representations. Such representations abstract away from the quantitative representations that underlie many physically situated systems, providing more compact, structured representations that omit (unnecessary) detail. Qualitative representations exist for many aspects of space and time; action; uncertainty; and categorical knowledge (ontologies).

Qualitative representations have many advantages, including naturally encoding semantics for many systems, being accessible to humans, providing smaller state spaces for learning and reasoning, and also being suitable for communication through natural language. These advantages have seen them

being increasingly used in intelligent physically grounded systems, from encoding spatial configurations of objects to modeling human behavior over time. This work is being done in many different places and across different subfields of AI such as knowledge representation and reasoning, planning, uncertainty, learning, and perception.

Work from all of these subfields was presented at the AAAI Spring Symposium on Qualitative Representations for Robots. We had sessions on the qualitative representation of the movement of mobile robots around humans; qualitative knowledge for planning and plan execution; qualitative representations to aid computer vision and scene understanding; and approaches for performing machine learning on qualitative representations for robotic tasks. Some of the key issues that arose across multiple contributions, and during the lively discussions in the breaks, were the challenges of projecting back into quantitative forms from qualitative representations in order to generate behavior in the real world; the challenges of combining multiple (perhaps redundant) qualitative calculi in a single system; computing distance measures between qualitative states; and the variety of ways movements can be encoded qualitatively.

Three fantastic invited speakers augmented the contributed papers and presentations. Jeff Siskind from Purdue University and Anthony Tony Cohn from University of Leeds both presented hugely impressive overviews of their research with qualitative representations, particularly their use for activity recognition and scene understanding from sensor data. A key issue in both their talks was the exploitation of context (for example, the role an object plays in an activity) to combat uncertainty from noisy observations. Matt Klenk from PARC also addressed the importance of context, but related to the understanding of spatial regions. A special mention should also go to Christian Dondrup, a graduate student at the University of Lincoln, UK, who communicated the key issues from our symposium in an engaging way for the plenary, while managing to keep quiet about the complexities of the qualitative trajectory calculus.

Nick Hawes served as the chair of this symposium, with additional help from an international committee of colleagues. The papers from the symposium were published as AAAI Press Technical Report SS-14-06.

Social Hacking and Cognitive Security on the Internet and New Media

The massive explosion of behavioral data made available by the advent of social media has empowered researchers to make significant advances in our understanding of the dynamics of large groups online. However, as this field of research expands, so do the opportunities multiply to use this understanding to forge powerful new techniques to shape the behavior and beliefs of people globally. These techniques can be tested and refined through the data-rich online spaces of platforms like Twitter and Facebook.

These techniques may be put to ethical or unethical ends. One might imagine using these techniques to encourage positive norms and stamp out perni-

cious misinformation online. Alternatively, one might use these technologies to spread disinformation or to break apart existing social bonds and erode trust.

Cognitive security is a term that examines this evolving frontier and suggests that in the future researchers, governments, social platforms, and private actors may be engaged in a continual arms race to influence — and protect from influence — large groups of users online.

Although cognitive security emerges from social engineering and discussions of social deception in the computer security space, it differs in a number of important respects. First, whereas the focus in computer security is on the influence of a few individuals, cognitive security focuses on the exploitation of cognitive biases in large public groups. Second, while computer security focuses on deception as a means of compromising computer systems, cognitive security focuses on social influence as an end unto itself. Finally, cognitive security emphasizes formality and quantitative measurement, distinct from the more qualitative discussions of social engineering in computer security.

Our AAAI Spring Symposium was titled Social Hacking and Cognitive Security on the Internet and New Media and attempted to survey the frontier of work being done in the space on a theoretical and practical level. Researchers Anil Vullikanti (Virginia Tech) and Huan Liu (Arizona State) examined the mathematical models used in describing the spread of misinformation through social networks. Researcher Paulo Shakarian (West Point) discussed the use of these network models in prototype software called GANG, which is currently in use by the Chicago Police Department, in targeting intervention efforts to persuade gang members to leave criminal activity.

One consistent thread through these discussions was the extent to which realistic bot identities on social media may become a powerful means through which these techniques are deployed. Swarms of bots may be used to create the illusion of an upswell of social activity or provide peer pressure to persuade unsuspecting groups of users online. Incidents in Syria, Russia, Turkey, and Mexico featuring large numbers of bots targeted at suppressing dissenting voices on social media suggest that this is already underway.

The emergence of bots as a tool of influence online will also probably be driven by the extensive financial incentives to produce and maintain farms of realistic personas online. Researcher Yazan Boshmaf (University of British Columbia) discussed the profit motivation in creating these identities and the discovery of large groups of these false personas on platforms like Facebook.

Insofar as bots appear poised to be one of the most significant real-world manifestations of the emerging research into cognitive security, our symposium produced a number of challenges to the artificial intelligence community; specifically: (1) Can campaigns of malicious bot influence be reliably detected? What patterns should be used to determine whether an account is a real person or a bot with an influence agenda? (2) How might bots be developed that learn effectively from their surroundings on social media? How might we measure their success online? (3) Can bots be designed to identify and resist misinformation on a mass scale online? How would one design agents to do this effectively? (4) What are the rules of ethics that should be designed around the deployment of bots for persuasive purposes online? (5) How might teams of bots be programmed to work together to achieve a social objective? How could this be modeled in an intelligent system?

Tim Hwang and Rand Waltzman served as cochairs of this symposium. No technical report for this symposium was published.

Manish Jain is an assistant professor in the Computer Science Department at Virginia Polytechnic Institute and State University.

Albert Xin Jiang is a postdoctoral researcher in the Computer Science Department at the University of Southern California.

Takashi Kido is a research manager at Rikengenesis in Japan. He had been a visiting researcher at Stanford University.

Keiki Takadama is a professor at the University of Electro-Communications in Japan.

Eric G. Mercer is an associate professor from the Computer Science Department at Brigham Young University, Provo, Utah.

Neha Rungta is a research scientist at NASA Ames Research Center, Moffet Field, California.

Mark Waser is CTO of the Digital Wisdom Institute.

Alan Wagner is a senior research scientist at the Georgia Tech Research Institute, Georgia Institute of Technology.

Jennifer Burke is a systems engineer in human factors at Boeing Research and Technology.

Don Sofge is a computer scientist at the Naval Research Laboratory.

William Lawless is a professor of mathematics and psychology at Pain College, Augusta, Georgia.

Mohan Sridharan is an assistant professor of computer science at Texas Tech University.

Nick Hawes is a senior lecturer in the School of Computer Science at the University of Birmingham.

Tim Hwang is chief scientist of the Pacific Social Architecting Corporation and a researcher at the Data and Society Research Institute in New York City.