

Reports on the 2015 AAAI Workshop Series

Stefano V. Albrecht, J. Christopher Beck, David L. Buckeridge, Adi Botea, Cornelia Caragea, Chi-hung Chi, Theodoros Damoulas, Bistra Dilkina, Eric Eaton, Pooyan Fazli, Sam Ganzfried, C. Lee Giles, Sébastien Guillet, Robert Holte, Frank Hutter, Thorsten Koch, Matteo Leonetti, Marius Lindauer, Marlos C. Machado, Yuri Malitsky, Gary Marcus, Sebastiaan Meijer, Francesca Rossi, Arash Shaban-Nejad, Sylvie Thiebaut, Manuela Veloso, Toby Walsh, Can Wang, Jie Zhang, Yu Zheng

■ *AAAI's 2015 Workshop Program was held Sunday and Monday, January 25–26, 2015, at the Hyatt Regency Austin Hotel in Austin, Texas, USA. The AAAI-15 workshop program included 16 workshops covering a wide range of topics in artificial intelligence. Most workshops were held on a single day. The titles of the workshops included Algorithm Configuration; Artificial Intelligence and Ethics; Artificial Intelligence Applied to Assistive Technologies and Smart Environments; Artificial Intelligence for Cities; Artificial Intelligence for Transportation: Advice, Interactivity, and Actor Modeling; Beyond the Turing Test; Computational Sustainability; Computer Poker and Imperfect Information; Incentive and Trust in E-Communities; Knowledge, Skill, and Behavior Transfer in Autonomous Robots; Learning for General Competency in Video Games; Multiagent Interaction without Prior Coordination; Planning, Search, and Optimization; Scholarly Big Data: AI Perspectives, Challenges, and Ideas; Trajectory-Based Behaviour Analytics; and World Wide Web and Public Health Intelligence.*

Algorithm Configuration

Automatic algorithm configuration addresses the problem of determining the settings of an algorithm's parameters to optimize its performance. It has most prominently been applied to optimize solvers for hard combinatorial problems (for example, SAT, MIP, ASP, and AI planning) as well as to hyperparameter optimization of flexible machine-learning frameworks (such as deep neural networks, or the space of algorithms defined by the Waikato environment for knowledge analysis, WEKA), but it also has applications in many more areas.

Additionally, algorithm configuration has been performed manually by domain experts in a tedious and time-consuming optimization process, a task humans are poorly suited for.

Nowadays, modern algorithm configuration procedures, such as ParamLLS, GGA, irace, SMAC, and ReACT, are used to perform this task automatically, saving valuable human time and allowing for a more systematic and reproducible way to determine well-performing parameter configurations. Thereby, automatic algorithm configurators often improve the performance of an algorithm by several orders of magnitude.

The formal task of algorithm configuration consists of determining a configuration c of a configuration space C from an algorithm A on an instance set I by optimizing a given performance metric $m : C \times I \rightarrow ?$ over all instances I . In contrast to the area of continuous black-box optimization (tackled, for example, in the evolutionary algorithms literature), algorithm configurators have to deal with several extensions: (1) different types of parameters (in particular, real-valued, integer-valued and categorical parameters, as well as conditional dependencies between parameters); (2) optimization of performance over a set of instances; and (3) highly stochastic and/or expensive evaluations of a configuration (for example, when this evaluation solves a SAT problem).

This was the first workshop specifically on algorithm configuration. However, it follows a line of related workshops and tutorials, for example, several workshops on configuration and selection of algorithms (COSEAL), the workshop on automated selection and tuning of algorithms at PPSN 2012, the Bayesian optimization workshops at NIPS 2012, NIPS 2014, the AutoML workshop at ICML-14, and the tutorial on algorithm selection and configuration at AAAI-13. What set this workshop apart is its specific focus on the algorithm configuration problem, and its aim to bring researchers together that develop and apply algorithm configuration methods to initiate dialogue, provide an empirical foundation for research in the field, discuss the respective strengths and weaknesses of different approaches, and determine promising directions for future research.

To promote these discussions, the workshop featured a mixture of invited talks, posters, presentations, and a lively panel discussion. While the posters and presentations displayed some of the modern applications of algorithm configuration, the invited speakers provided a bigger picture. Meinolf Sellmann presented his view on algorithm configuration influenced by real-world experience at IBM, Kevin Leyton-Brown (University of British Columbia) discussed the programming by optimization paradigm that is enabled through the existence of effective configurators, and Franco Mascia (Université Libre de Bruxelles) showed how algorithm configuration enables the grammar-based construction of algorithms. The closing panel discussion allowed for an open discussion among the active researchers in the area to express their views on the future of the field. A major

topic of discussion was the necessity of easy-to-use general algorithm configuration methods that lower the barrier of entry for researchers outside of the community (so as to increase visibility). Another prominent topic of discussion was the opportunity to explore additional application domains, especially in regards to big data.

Automated algorithm configuration is a powerful tool that enables the automatic customization of complex (and flexible) solvers to achieve peak performance for specific benchmarks of interest, while at the same time reducing the need for manual work by domain experts. Modern techniques are now able to offer substantial performance improvement for solvers with tens to hundreds of parameters, and as a consequence, algorithm configuration is becoming ever more popular and is taking off in new and exciting directions.

Frank Hutter, Marius Lindauer, and Yuri Malitsky wrote this report and cochaired the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-15-01.

Artificial Intelligence and Ethics

Concern about progress in AI is a common theme of the zeitgeist today. Famous scientists like Steven Hawking and technologists like Elon Musk have frequently commented on their fears about where AI may take us. To begin a discussion around these and other issues within the community of AI researchers, several events were held at AAAI-15 including a debate on the proposed UN ban on lethal autonomous weapons and this workshop. Topics covered at the workshop included architectures for endowing autonomous agents with ethics, prospects for the future of AI, and experiments around AI and ethics.

One highlight of the workshop was the invited talk by Steve Goose, director of the Human Rights Watch Arms Division. He gave a talk on the progress toward a UN ban on lethal autonomous weapons. This provided background detail to the debate held within the main conference with Ron Arkin, associate dean for research and space planning at Georgia Tech. Goose was instrumental in bringing about the 2008 convention banning cluster munitions, the 1997 treaty banning antipersonnel mines, the 1995 protocol banning blinding lasers, and the 2003 protocol requiring cleanup of explosive remnants of war. In 2013, he and Human Rights Watch cofounded the Campaign to Stop Killer Robots. Goose was thus well placed to speak at length about these issues. We are very grateful to him for finding the time to talk to the workshop.

The workshop also included a wonderful video (available on YouTube) by Herb Simon on Forecasting the Future or Shaping It? Manuela Veloso introduced the video, putting it in context and inspiring us all to follow up on Herb Simon's ideas and suggestions. You

are encouraged to go listen to the video if you haven't already seen it. The workshop also united with the sister workshop on Beyond the Turing Test with a concluding panel discussion on the future of AI.

Discussions are now underway to have a special issue around selected presentations from the workshop. In addition, given the good response to the workshop, we intend to run a followup event at AAAI-16, with a preliminary title of AI, Ethics, and Society. In addition to the theme discussed so far, there will also be an increased emphasis on the impact of AI on society (for example, its impact on the workplace).

Toby Walsh chaired this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-15-02.

Artificial Intelligence Applied to Assistive Technologies and Smart Environments

This 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 presented at the workshop was ambient assisted living (AAL), which included six papers. A first talk given by Sébastien Guillet outlined general principles to model and synthesize the control part of AI systems that can adapt and optimize the behavior of a smart home system to prevent potentially harmful situations. Julien Maitre gave two talks; the first presented an improved methodology for activity recognition using electrical devices identification, and the second presented a complementary recognition system that is based on environmental sounds and that aims to detect errors related to cognitive impairment. Then a talk given by Thomas E. Allen discussed the preferences of the user in smart homes and assistive environments and the algorithmic ways to take them into account in the assistance process. A talk given by Aarti Malhotra presented a user study that aims to understand how audiovisual prompts are perceived by humans on an emotional level. And finally, a talk given by Boris Galitsky presented an evaluation of the assistance with reasoning about the mental world for children with autism.

Another major theme was human-robot interaction. The papers on this theme described robotic systems and frameworks dedicated to user assistance. The papers of Chung Hyuk Park and Sima Behpour presented robotic frameworks respectively dedicated to interaction with children with autism and learning tasks. A talk given by Bonnie Dorr presented an artificial companion for daily help to patients who

suffer from traumatic brain injury (TBI) or Amyotrophic Lateral Sclerosis (ALS). David Miller gave a talk about a robotic support system that focused on kinematic capture and EEG monitoring and targeting children with cerebral palsy to help them develop their crawling abilities. Finally, Tiaro Vaquero presented the implementation of a system dedicated to the planning and scheduling of activities of daily living, involving multiple users (especially elders) and human-robot interaction with multiple mobile assistive and collaborative robots.

The third and last theme of the workshop was the use of speech analysis and recognition, where, first, Bonnie Dorr proposed a future investigation position on a speech analysis tool that follows the progression of ALS disease and recognizes patient speech to increase the patient's autonomy in a smart home; then Bonny Banerjee presented an online unsupervised algorithm for learning the speech of severely to profoundly hearing impaired patients.

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.

Bruno Bouchard served as chair of this workshop. Sylvain Giroux, Abdenour Bouzouane, and Sébastien Guillet served as cochairs. This report was written by Sébastien Guillet. The papers of this workshop were published as AAAI Press technical Report WS-15-03.

Artificial Intelligence for Cities

Almost half of humanity today lives in urban environments and that number will grow to 80 percent or more by the middle of this century in different parts of the world. Cities are thus the loci of resource consumption, economic activity, social interactions, and education and innovation; they are the cause of our looming sustainability problems but also where those problems must be solved. Cities are also an enormous forum for policy making, as well as an apparently unbounded source of digital data of a wide nature. Artificial intelligence has the potential to play a central role in tackling the underlying hard computational, decision-making, and statistical challenges of smart cities.

The workshop brought together leading researchers from academe, industry, and city agencies, working on a wide range of AI and data science subfields such as machine learning, optimization,

visualization, database systems, knowledge representation, and planning. The first part of the workshop included invited talks from experts across these fields. Juliana Freire at New York University spoke about how big data techniques can be married with interactive visualization and analysis to enable users to find patterns in urban activity such as taxi data. Manuela Veloso at Carnegie Mellon University spoke on optimizing traffic flow with artificial intelligence techniques in cities and robot localization (SLAM) in urban environments. Craig Knoblock at the University of Southern California spoke on how differing ontologies can impede cross-city data analysis and how to automate the detection and merging of public data-set ontologies. Adi Botea at IBM discussed edging risk in journey planning with a risk-averse multimodal journey advisor that performs journey planning, journey monitoring, and replanning. Finally, Mike Flowers at Enigma and New York University discussed data science in New York City, including fire-risk models, analyzing NYC business locations and implementing DataBridge.

The workshop concluded with contributions in the areas of building energy efficiency, predictive policing in New York city, trajectory tracking, optimizing bike-share distribution, and urban sensing. The workshop participants discussed data access in cities, data generating processes and inherent biases in the collection protocols, challenges associated with open city data, acquiring domain knowledge and expertise, impacting and informing policy, and finally, the need for a tighter integration of the AI community and its subfields to make an impact on urban science and city governance.

AI research has the opportunity to transform our cities around the world for the better by helping improve operations, services, security, citizen participation, and quality of life.

Theo Damoulas wrote this report and served as chair of the event. The papers of the workshop were published as AAAI Press Technical Report WS-15-04.

Artificial Intelligence for Transportation: Advice, Interactivity, and Actor Modeling

The transportation domain is increasingly taking up AI techniques in the design of products and systems. Cars nowadays implement machine-learning algorithms. When searching for a route on a mobile app, solutions are provided through AI algorithms. This workshop has covered a diverse selection of topics, both from a more theoretical and more practical nature, showing once again that there are many transportation problems where applying AI technology is beneficial. Traffic signal control was featured in two presentations, one of which was an invited talk given by Scott Sanner, a principal researcher at NICTA and the Australian National University.

The problem of fairly distributing the costs to a set of clients is important, among other domains, in vehicle-routing problems, where a fleet of vehicles deliver goods to a set of clients. While Shapley cost allocations are known to be optimal when a few relevant assumptions are imposed, the exact computation of such optimal allocations is a very challenging computational problem. One presentation focused on new hardness results, as well as a number of more tractable, but not necessarily optimal techniques.

The increasing availability of car parking data makes it possible to utilize AI algorithms for a more effective use of car parking lots. More specific topics covered in the workshop include pricing strategies aimed at maximizing the occupancy of parking lots, and predicting the occupancy of parking spots within a time window in the future.

In an electric vehicle, the amount of the battery power available is a valuable resource. Battery power is needed not only to power the engine, but also to run auxiliary systems, such as the climate-control system. Part of the work presented in the workshop has focused on developing an adaptive-advised agent that makes recommendations to the driver about the settings of the climate-control system.

Navigating in a hostile environment brings up the need to avoid, as much as possible, potential ambush locations set by an adversary. This was the topic of one presentation, which presented strategies for planning routes optimized for ambush avoidance.

Aerial transportation was another well-represented topic. One presentation addressed the problem of planning the trajectory of a helicopter or tilt rotor craft while respecting a number of noise-related constraints. Another presentation focused on towing aircraft at an airport with self-driving vehicles.

Most work featured in the workshop was motivated by important real-world problems, evidence of the important maturation of the community of researchers working on the intersection of AI and transportation.

The last part of the workshop was an open discussion that revisited at a deeper level the topics presented and discussed earlier in the day. At the conclusion of the event, there was a consensus that similar events should continue to be organized, possibly in colocation with major AI conferences.

Adi Botea and Sebastiaan Meijer wrote this report and served as co-chairs of the event. The papers presented in the workshop are available as AAAI Technical Report WS-15-05.

Beyond the Turing Test

The Turing test, now more than 60 years old, has long served as a highly visible, public signpost for research in artificial intelligence. But competitors like Eugene Goostman and PARRY often seem like exercises in evasion, rather than robust advances in

general intelligence; after 60 years, the test itself might be due for a refresh. Inspired in part by an article by Gary Marcus in the *New Yorker*, we have become engaged in an effort to go “Beyond the Turing Test.”

The purpose of the workshop, modeled on a set of early meetings that helped shape the annual RoboCup competitions, was to seek community input. More precisely, the goal was to craft a replacement, an annual or biannual Turing championship, that might consist of three to five different challenging tasks, with bragging rights given to the first programs to achieve human-level performance in each task.

The workshop was organized with 12 invited talks and 13 posters, with all the participants showing a lot of enthusiasm for the initiative. The presentations and group discussions converged to possible tests and challenges along four main lines. (1) *Commonsense knowledge understanding*: The Winograd Schema Challenge, recently sponsored by Nuance and proposed by Hector Levesque, tests the ability to resolve linguistic antecedents in contexts in which commonsense knowledge is critical. (2) *Integrated language and image question answering*: The focus is on the comprehension of novel materials, such as videos, texts, photos, and podcasts. (3) *Task-based perception and actuation*: The goal is to go beyond language and address the embodiment of intelligence in a task involving the understanding of the physical qualities, such as assembling a piece of furniture, possibly in collaboration with a human. (4) *Discipline-specific knowledge testing*: From a cognitive science point of view, we would also be interested in matching the mechanisms of human performance in discipline-specific tests, such as fifth-grade math or physics.

We also discussed how to design an AI competition for an inducement prize, also by taking inspiration from the RoboCup approach. The people involved are ready to work on the definition and implementation of the new tests. A follow-up workshop will take place at IJCAI 2015 in Buenos Aires, Argentina. Our plan is to have the first new tests in place for AAAI 2016.

Gary Marcus, Francesca Rossi, and Manuela Veloso served as cochairs of the workshop and wrote this report. No technical report was published.

Computational Sustainability

Computational sustainability is a fast-growing interdisciplinary field that applies techniques from computer science, information science, operations research, applied mathematics, and statistics to solve problems that involve balancing environmental, economic, and societal needs for sustainable development. Computational sustainability spans a multitude of sustainability-related topics such as biodiversity conservation, energy, urban planning,

climate change, transportation, water, food, health, and poverty.

Importantly, the computational problems that arise in many sustainability domains relate to a wide spectrum of AI topics and techniques, such as graphical models and probabilistic inference, statistical learning, data and graph mining, constrained and stochastic optimization, reasoning under uncertainty, spatiotemporal modeling, and network science.

The Association for the Advancement of Artificial Intelligence has played a key role in fostering the growth of this new research field by organizing a special track on computational sustainability as part of the main AAAI conference since 2011. The collocation of this workshop with the 2015 AAAI conference provided a more interactive and informal forum, where researchers interested in this field could present preliminary work, hear about key topics from invited speakers, get to know each other, and engage in discussions.

The workshop attracted more than 50 participants from diverse backgrounds, institutions, and countries. It included 31 papers, a combination of position and full papers. Two topics were particularly well represented in the accepted papers: energy and environmental applications. In the energy domain, the papers touched on forecasting energy demand, incorporating renewables, dynamic demand response, energy disaggregation, leveraging smart meter data, autonomous electricity trading, and sustainable building design, among others. Related to the environment, several papers addressed issues in ecology such as dynamic reserve design, inferring migratory routes, identifying biodiversity hotspots, acoustic species identification, game-theoretic aspects of wildlife and fisheries protection, and assessing ecological integrity. Others looked at land-cover classification, climate negotiations, materials discovery, sustainable supply chains, water networks, and epidemic spread.

The workshop also included four invited talks spanning various subjects across the sustainability and computational spectrum. Claire Monteleoni (George Washington University) spoke about climate informatics and the role that machine learning can play in improving predictions of climate-change trends and extreme events. Milind Tambe (University of Southern California) highlighted the need to extend game-theoretic ideas successfully harnessed by security agencies to protect airports, the coastal waters, trains, ferries, and campuses to design green security games, focused on deploying limited security resources to protect forests, fish, and wildlife. In the context of facilitating the use of alternative and sustainable forms of transportation, Andreas Krause (ETH Zurich) spoke about a crowdsourcing mechanism that incentivizes the users of a bike-sharing system to help in the costly bike-repositioning process. Finally, Manish Marwah (HP Lab) gave an overview

on the use of data analytic techniques for the environmental footprinting of products. At the end of the workshop, the invited speakers were joined by Carla Gomes (Cornell), one of the founders of the field of computational sustainability, in a panel discussion on the directions and challenges for computational sustainability. The panelists and workshop attendees exchanged ideas about the future of the field and on strategies for growing a productive and successful research community.

Bistra Dilkina (Georgia Tech) was a coorganizer and author of this report. Stefano Ermon (Stanford University), Rebecca Hutchinson (Oregon State University) and Daniel Sheldon (University of Massachusetts, Amherst) were also organizers of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-15-06.

Computer Poker and Imperfect Information

Poker has emerged as a major AI challenge problem. Since 2006, there has been a competition between the strongest computer poker agents held annually at AAAI. Building strong poker agents involves solving many of the most challenging issues faced by all problems in artificial intelligence, including dealing with enormous state spaces, the presence of multiple self-interested agents, imperfect information, stochastic events, balancing between exploitation of opponents' mistakes and minimizing our own exploitability (ensuring that strong agents cannot take advantage of our own mistakes), and performing robust, large-scale optimization. Poker is not simply a toy game; it is tremendously popular for humans, and online poker is a multibillion dollar industry. The version of two-player no-limit Texas Hold'em played in the AAAI competition has approximately 10^{165} states in its game tree. Other domains where algorithms for imperfect-information games are pivotal include business (for example, auctions and negotiations), medicine (for example, developing robust treatment policies to combat diseases), and (cyber)security.

The emergence of poker as a major challenge brought together researchers studying a variety of topics pertaining to imperfect-information games, ranging from theoretical analysis and general-purpose algorithms to the design of agents for poker and other domains. The workshop was conceived at AAAI in 2012, and has now run for four consecutive years. The 2015 workshop included nine oral presentations followed by a poster session and round table discussion. The standard paradigm used by most strong agents for large imperfect-information games is first to apply an abstraction algorithm to compute a smaller approximation of the game, and then to compute an approximate equilibrium in the abstraction; these are done offline, and then the strategies

are implemented in real time by performing a table lookup. Several of the talks proposed novel approaches constituting significant departures from this paradigm. Kevin Waugh from Carnegie Mellon University described a new game-solving approach called functional regret estimation, which combines the abstraction and equilibrium-finding components by employing an online regressor from domain features to approximate regrets used by a regret-minimizing equilibrium-finding algorithm. This, in effect, generalizes the standard abstraction approaches by allowing us to learn good abstractions from features, as opposed to computing a single fixed abstraction in advance of equilibrium finding. For the application of the approach to poker, these features included expected hand strength, pot size, and number of actions in the current hand.

Sam Ganzfried from Carnegie Mellon University described another approach that marks a significant departure from the traditional paradigm, involving solving portions of the game that are actually reached in real time to a greater degree of accuracy than is done in offline computations. A new efficient algorithm for performing such endgame solving in large imperfect-information games was presented, and the approach was shown to lead to significant performance improvements in no-limit Texas Hold'em against the strongest competition agents. Theoretical analysis showed that this approach may produce highly exploitable strategies in certain games, although it guarantees low exploitability in others, and a new framework was presented for assessing its broader applicability in different game classes.

As is typical for this workshop, most of the talks presented novel theoretical analysis and domain-independent approaches, and several domains were considered in addition to poker. Hsuan-Tien Lin from National Taiwan University presented an approach for learning new bidding conventions in the domain of bridge. Training on random deals, their algorithm was able to learn a bidding system that performed competitively with a champion bridge agent that uses the Standard American human bidding system. The conventions learned by the algorithm differed significantly from the standard system; for example, for its initial action the new approach bids one spade with a hand with nearly balanced suits and never bids one notrump, while the standard approach bids one spade with five or more spades, and frequently bids one notrump.

Kevin Waugh from Carnegie Mellon University presented the results from the poker competition. Due to the close proximity to the most recent competition, which took place in July 2014, this competition was significantly scaled down and featured only three-player Kuhn poker. A full-scale competition is planned for the 2016 AAAI conference. Four agents were submitted to the 2015 three-player Kuhn poker competition; the winning agent was Umpa,

from Ariel University. The agent used an opponent modeling approach where each opponent was represented by a vector of 48 real numbers in $[0,1]$, corresponding to each of the possible 48 states of the game (where each state consists of a private card and betting history for the current round). The models were updated based on new information revealed after each hand, utilizing a genetic algorithm to tune the learning rate.

The round table discussion focused on possible rule changes and integration of new events for the 2016 poker competition. There was a strong contingency in favor of adding a six-player no-limit Texas Hold'em competition. Some of the arguments in favor were that it is an extremely popular variant for humans, that it would involve drastically new algorithmic challenges to deal with the massively larger state space, and that it would involve conceptual challenges, as Nash equilibrium is not as well motivated as in two-player games. Some arguments against this proposal were the increased computational resources required, as well as the challenges with obtaining statistical significance. However, such challenges also open up interesting research questions on tournament design and variance reduction.

There was also interest in adding a lightweight two-player no-limit Texas Hold'em competition, with a significantly smaller limit on the size of agents. This would potentially make the competition more easily accessible to newer entrants and approaches that may have more limited access to computational resources than some existing teams. Kevin Waugh will be one of the organizers for the 2016 competition, and the other organizer is still to be determined.

Sam Ganzfried served as the chair of this workshop and wrote this report. The papers of the workshop will be published as AAAI Technical Report WS-15-07.

Incentive and Trust in E-Communities

The area of trust has experienced rapid growth in the past decade. With the growing prevalence of social interactions through electronic means, trust becomes considerably important. Many computational and theoretical models of trust that are well-suited for a variety of domains, such as e-commerce, social networks, web blogs, ad hoc networks, and others, have been recently developed. They present trust as a multifaceted concept that operates at many levels and plays important roles in ensuring reliable interactions. Although trust-enabled systems allow people to act under uncertainty and mitigate the risk of negative consequences, still, sociotechnical attacks often succeed by exploiting loopholes in the design of trust systems. 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 game-theoretic mechanisms have been developed to discourage deception and fraud in e-communities, promote honest behaviors, and create incentive for participants to contribute truthfully; that is the nice property of incentive compatibility.

Trust and incentive have bidirectional relationships. As trust 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 systems should not only provide a means to detect and prevent malicious activities but also design a mechanism to discourage dishonest attitudes among participants. The evidential success of combining these two concepts inspires and encourages researchers in the trust community to enhance the efficacy and performance of trust-modeling approaches by adopting various incentive mechanisms. The main 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 modeling, toward the design of more effective trust and incentive mechanisms for creating safe electronic communities.

There were three main themes of papers in this workshop: (1) the design of effective trust models; (2) the design of incentive mechanisms from the game-theoretic perspective; and (3) the combination of both. Papers concerned with the first theme discussed trust models that assist users that achieve high levels of trust to be selected as partners of others in e-communities for future interactions. For example, the paper by Abdullah Aref and Thomas Tran presented a decentralized model for establishing trust in multiagent systems where trustees predict the demands of the trusters, modifying their behavior to satisfy the need of trusters. Another paper, by Boris Galitsky, presented a conversational agent performing social promotion (CASP) to assist in automation of interacting with Facebook friends and managing other social network contacts. A third paper, written by Mehrdad Nojournian, tried to understand how humans gain or lose trust in their daily life interactions and how behavior and attitudes of humans can be controlled in various social encounters. Papers concerned with the second theme presented strategy-proof mechanisms where agents are better off telling the truth about their private information, such as needs, values of goods, and so on. The paper of Ferran Torrent and colleagues, which dealt with the third theme, discussed a method to manage trust in multiattribute auctions to achieve incentive compatibility.

The workshop participants discussed the various angles where the two areas (trust and incentive) can be merged, what benefit can be brought by doing so, and what additional efforts the researchers in these

two areas can put to advance this emerging direction.

Jie Zhang, Zeinab Noorian, Stephen Marsh, and Christian Jensen served as cochairs of this workshop. Jie Zhang wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-15-08.

Knowledge, Skill, and Behavior Transfer in Autonomous Robots

Autonomous robots have achieved high levels of performance and reliability at specific tasks. However, for them to be practical and effective at everyday tasks in our homes and offices, they must be able to learn to perform different tasks over time, demonstrating versatility. Learning each task in isolation is an expensive process that requires large amounts of both time and data. In robotics, this expensive learning process also has secondary costs, such as energy usage and joint fatigue. Recent developments in transfer and multitask learning provide a potential solution to this problem, enabling robots to minimize the time and cost of learning new tasks by building upon knowledge learned from other tasks. This ability is essential to enable the development of versatile autonomous robots that are expected to perform a wide variety of tasks and rapidly learn new abilities.

Various aspects of this problem have been addressed by research across several different communities, including machine learning, knowledge representation, optimal control, and robotics. This workshop brought together researchers from these different communities toward the goal of enabling autonomous robots to support a wide variety of tasks, rapidly and robustly learn new abilities, adapt quickly to changing contexts, and collaborate effectively with other robots and humans to achieve a common goal. One of the main themes of the workshop was transfer learning in a reinforcement learning setting. Automatic hierarchical decomposition of tasks and policies was explored, together with abstract representations that can capture the similarities between tasks. Matthew Taylor (Washington State University) gave an invited talk on different aspects of transfer learning, and in particular between different agents that can communicate only through advice.

Another theme was knowledge transfer between humans and robots. Robots can positively affect human behavior, especially with children, who can learn from the interaction with robots differently from their interactions with humans. Furthermore, robots can exploit interaction in order to change their own behavior and to acquire knowledge about objects and categories of interest, in particular when encompassing different sensory modalities. Workshop participants also discussed environments and benchmarks for knowledge transfer in robotics. Maria Gini (University of Minnesota), in her invited

talk, presented different complex scenarios for multi-robot systems, including the RoboCup rescue simulator. The workshop participants discussed with interest the diversity of the methods in this emerging area and the difficulties that still arise in their application to physical robots. The challenge of knowledge transfer can take different shapes in the different fields related to artificial intelligence and robotics. Participants agreed that future workshops on this topic will help identify connections across such fields, in order to overcome the difficult problem of overly specialized robots, unable to generalize to similar contexts and exhibit versatile behaviors.

Matteo Leonetti served as chair of the workshop, with the collaboration of Eric Eaton and Pooyan Fazli as cochairs. The papers of the workshop were published as AAAI Press Technical Report WS-15-09.

Learning for General Competency in Video Games

One of the longstanding goals of AI research is the design of generally competent agents: agents capable of achieving strong performance in many, varied, interesting, and challenging domains. In order to evaluate general competency, an emerging trend in AI research has been to turn to video games and video game platforms. The Atari 2600, for example, offers hundreds of independently designed games drawn from a variety of genres. It satisfies our needs for diversity and interestingness, and challenges existing general-competency methods and AI techniques. As a whole, video game platforms offer a broad spectrum of research opportunities, including model learning, planning, transfer learning, apprenticeship and imitation learning, intrinsic motivation, and representation learning.

With the recent release of the arcade learning environment (ALE), which provides an easy-to-use reinforcement learning interface to over 55 Atari 2600 games, a number of groups from around the world have initiated Atari 2600-related projects. Although the platform has led to some high-profile success stories, including the much-publicized Deep Q-Networks of Volodymyr Mnih and colleagues, it still offers many unsolved challenges. For example, learning dynamical models for high-dimensional video input still remains unsolved, as does learning concise representations amenable to planning algorithms. In addition, effective exploration when rewards are very sparse is proving to be a central challenge behind many of the harder Atari 2600 games.

This workshop brought together researchers working on video game-related research to further stimulate the development of learning methodologies suited to assess general agent competency. Its sessions were planned to enable the exposure of challenges and current state-of-the-art approaches, as well as to enable discussions on how best to

strengthen and improve existing evaluation methodologies.

The first part of the workshop revolved around new algorithms directly tackling some of the AI challenges posed by video game platforms. Michael Bowling (University of Alberta) listed several challenges that can be addressed using the ALE as an evaluation platform. Erik Talvitie (Franklin and Marshall College) and Michael Bowling proposed a new simple feature set for reinforcement learning in visual domains, designed to capture pairwise, position-invariant, spatial relationships between objects. Marlos Machado (University of Alberta) and colleagues presented a domain-independent optimistic initialization approach for reinforcement learning. Satinder Singh (University of Michigan) presented some of his recent work, showing, for example, how one can generate a real-time player from planning and deep-learning techniques. Nir Lipovetzky (University of Melbourne) and colleagues discussed how one can use classical planning algorithms without having a PDDL-model nor any prior knowledge of the actions effects and goals. Matthew Hausknecht (University of Texas at Austin) concluded this portion of the workshop with some recent results applying neuroevolution to policy search for the Atari 2600.

The second half of our workshop focused on the evaluation of generally competent agents. Peter Stone (University of Texas at Austin) recapitulated some of the important evaluation lessons learned from the General Game Playing competition. Matthew Hausknecht and Peter Stone further spoke of the dangers of deterministic evaluation. In addition, Marc G. Bellemare (Google DeepMind) provided empirical evidence of the exploitability of determinism. He presented an algorithm, the Brute, which optimizes a single game trajectory using an open-loop control approach. This led to two rounds of panel discussions where we agreed that deterministic evaluation takes us away from the goals of reinforcement learning. Also, the workshop participants came up with a set of evaluation standards to be followed, such as the essential information to be reported in future works. Moreover, we discussed the best way to inject stochasticity in the ALE. The panel was undoubtedly successful and led to the drafting of a set of evaluation standards for general competency in video games. We expect these standards to ease reproducibility and comparability between different research groups.

We believe the workshop was very successful, achieving all of its original goals. The organizers are now writing an article to present the evaluation standards discussed in the workshop. This article will also introduce a revised arcade learning environment, which will facilitate the new evaluation standards agreed upon during the workshop.

Marc G. Bellemare (Google DeepMind), Michael Bowling (University of Alberta, Canada), Marlos C.

Machado (University of Alberta, Canada), Erik Talvitie (Franklin and Marshall College, USA) and Joel Veness (Google DeepMind) organized this workshop. This summary was written by Marlos C. Machado. The papers of the workshop were published as AAAI Press Technical Report WS-15-10.

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 and rewards, and symbolic methods often make strong assumptions regarding norms and protocols. However, in realistic problems, these assumptions are easily violated. Thus, there is a need for new models and algorithms that specifically address the case of ad hoc interactions.

The purpose of this workshop was to discuss the role of such predefined knowledge and coordination in multiagent systems, and to provide a venue for research on novel models and algorithms that specifically address multiagent interaction without prior coordination (MIPC). There were a total of seven accepted papers, with topics as diverse as nonparametric Bayesian learning in I-POMDPs, optimal selection of multirobot coalition formation algorithms, combining the expert and type-methodologies for effective interaction, and the RoboCup 2014 SPL drop-in player competition. 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 again privileged to have invited talks by three distinguished researchers: “Leveraging Expert Feedback in Recommender Systems” by Pascal Poupart from the University of Waterloo; “Agent-Human Interaction without Prior Communication” by Sarit Kraus from Bar-Ilan University; and “Interactive POMDPs” by Piotr Gmytrasiewicz from the University of Illinois at Chicago.

The workshop was chaired by Stefano Albrecht, Jacob Crandall, and Somchaya Liemhetcharat. Stefano Albrecht was the author of this report. The advisory committee consisted of Subramanian Ramamoorthy, Peter Stone, and Manuela Veloso. The chairs would like to thank the workshop participants, the invited speakers, the program committee,

and the AAAI staff for making the workshop a success. A follow-up workshop is being planned for AAAI-16. The papers of the workshop were published as AAAI Press Technical Report WS-15-11.

Planning, Search, and Optimization

To facilitate mutual understanding across disciplinary boundaries the workshop included two one-hour tutorials: An Introduction to Search and OR in AI Planning, by Patrik Haslum and Menkes van den Briel; and Search in Integer Programming, by Thorsten Koch. Both were highly informative and gave attendees the background knowledge, basic terminology, and overview of the state of the art needed to appreciate the papers presented in the workshop.

Fourteen papers were presented at the workshop. As hoped for, they covered a wide range of topics with some having a strong OR focus (for example, “Effect of Bundle Method in Distributed Lagrangian Relaxation Protocol”), some a strong planning / search focus (for example, “Enumerating Preferred Solutions to Conditional Simple Temporal Networks Quickly Using Bounding Conflicts”). Many of the systems described in these papers were hybrids containing an OR component (typically a solver for linear programs) and a planning search component. For example, the paper “Linear Programming for Heuristics in Optimal Planning” summarized a recent line of work in the planning community in which the heuristic value for a state is computed by solving a linear program. As a second example, “Approximate Uni-Directional Benders Decomposition” describes how planning technology can be integrated into a general decomposition framework inspired by logic-based Benders decomposition. A number of papers additionally built on approaches from the formal methods community and integrated model checking with planning, search, or optimization. For instance, the paper “SMT-Based Nonlinear PDDL+ Planning” solves planning with nonlinear continuous change using a satisfiability modulo theory solver enriched with a new planning-specific variable selection heuristic.

Several papers in the workshop presented challenging real-world applications in which a hybrid OR-planning/search system was thought to be necessary for progress. Examples of such applications are seen in these papers’ titles: “A Realistic Multi-Modal Cargo Routing Benchmark” and “Preventing HIV Spread in Homeless Populations Using PSINET.”

The workshop provided very interesting insights into the similarity and differences between the methods and challenges in planning (AI) and optimization (OR). OR participants commented that they were pleasantly surprised by the interest of the AI community in OR methods. Following the workshop, we had several interesting discussions about possibilities to

build hybrid approaches. Currently planning is underway for visits of postdocs and more focused international workshops to investigate both the use of mathematical programming in planning and that of planning methods in optimization. While over time there has always been some interaction between these fields, during the workshop we developed a strong feeling that this should be increased.

J. Christopher Beck, Robert Holte, Thorsten Koch, and Sylvie Thiebaux were the authors of this report and also the cochairs that organized the workshop. The papers of the workshop were published as AAAI Press Technical Report WS-15-12.

Scholarly Big Data: AI Perspectives, Challenges, and Ideas

Academics and researchers worldwide continue to produce large numbers of scholarly documents, including papers, books, technical reports, and the like, and associated data, such as tutorials, proposals, and course materials. For example, PubMed has more than 20 million documents, 10 million unique names, and 70 million name mentions. Google Scholar has many millions more, it is believed. Understanding how and at what scale research topics emerge, evolve, or disappear, what is a good measure of quality of published works, what are the most promising areas of research, how authors connect and influence each other, who are the experts in a field, and who funds a particular research topic are some of the major foci of the rapidly emerging field of scholarly big data. The primary goals and objectives of the workshop were to promote both theoretical results and practical applications for scholarly big data, and address challenges faced by today’s researchers, decision makers, and funding agencies as well as well-known technological companies such as Microsoft and Google, repositories, and publishers such as Elsevier.

Papers presented at the workshop covered topics including identifying important citations in scholarly literature, that is, citations that indicate that the cited work is used or extended in new ways; identifying and extracting figures, tables, and captions from scholarly articles; learning multiple networks for author personalization and recommendation; analyzing research publications to understand gender preferences in computer science; automatically or semiautomatically acquiring, integrating, and using complex mechanism models, specifically related to cancer mechanisms, through automated reading and a hyperdetailed refinement process based on logical representations and powerful inference mechanisms; and examining the notion of the Big Data Lake and contrasting it with decision-support-based data warehouses.

The workshop also included two invited talks. The first talk, given by Yang Song (Microsoft Research

Redmond), focused on knowledge-powered next-generation scholarly search and recommendation engines that bridge the gap between unstructured text and structured relationships by converting the massive amount of unstructured scholar data on the web into structured entities in knowledge bases. The second talk, given by Vu Ha (Allen Institute for Artificial Intelligence, Seattle), focused on the semantic scholar: a search and discovery service for scientific papers with semantic capabilities.

Cornelia Caragea and C. Lee Giles served as cochairs of this workshop and coauthored this report. The papers of the workshop were published as AAAI Press Technical Report WS-15-13.

Trajectory-Based Behavior Analytics

In recent years, the data-driven scientific discovery approach has become an important emerging paradigm for computing in areas including social networks, services, cloud technology, and the internet of things. Under this paradigm, big data is the core that drives new research in many domains, from environmental to social. One important source of information for potential value creation is the real-time trajectory data obtained from entities including animals, robots, and humans. The trajectory information naturally reveals the details of instantaneous behaviors conducted by entities, which is closely related with complex behaviors in the form of multiple interdependent multivariate time series data with varied locations. This forms the need and emergence of behavior modeling (that is, understanding behaviors from cognitive and analytic perspectives) and behavior system construction (that is, developing cognition-as-a-service systems to support decision making).

The Trajectory-Based Behavior Analytics workshop focused on addressing deep science and research questions related to behavioral analytics for real-time trajectory-driven data applications as well as value delivery platform systems. This workshop included eight oral presentations, which were from the eight papers accepted on the basis of a rigorous peer-review process. The workshop brought together researchers from a variety of subfields of AI such as data mining, machine learning, and social networks.

Three of the presented papers at the workshop focused on pattern mining and learning applied to trajectory data. The authors proposed new algorithms to improve the current trajectory clustering accuracy, pattern-recognition precision, and active learning effectiveness. Two other presented papers were mainly about the visualization of trajectory behavior by using spatiotemporal relations and neural networks. One of them introduced a prediction model for the trajectory data through observation of prior behaviours. In addition, two other papers paid more attention to the application aspects in social networks when trajectory data is involved.

All the speakers in this workshop were well prepared and gave informative presentations, which attracted plenty of follow-up discussions. The workshop participants discussed how to represent and model trajectory behavior, how to incorporate context information into trajectory data, how to formalize a trajectory-based behavior network, how to integrate the multiple and heterogeneous trajectory sequences, and how to model the dynamic evolutions of trajectory-based behavior. All the participants took part in communicating various knowledge and information about the trajectory behaviour with other peers, and they would like to attend future workshops with topics and focus similar to this one.

Can Wang, Chi-Hung Chi, and Yu Zheng served as cochairs of this workshop and coauthored this report. This workshop has successfully provided a platform for research peers on trajectory studies to share knowledge and learn from one another. The papers of the symposium were published as AAAI Press Technical Report WS-15-14.

World Wide Web and Public Health Intelligence

Public health intelligence includes a set of activities to extract, capture, and analyze multidimensional socioeconomic, behavioral, environmental, and health data to support decision making to improve the health of the population. Advances in artificial intelligence tools and techniques and Internet technologies are dramatically changing the ways that scientists collect data and how people interact with each other and with their environment. Moreover, the Internet is increasingly used to collect, analyze, and monitor health-related reports and activities and to facilitate health-promotion programs and preventive interventions.

This workshop follows a successful first workshop held in July 2014 in Québec City, Canada, in conjunction with the 28th AAAI Conference on Artificial Intelligence (AAAI-14). The workshop brought together computer scientists, biomedical and health informaticians, researchers, students, industry professionals, and representatives of national and international public health agencies. Participants were interested in the theory and practice of computational models of web-based public health intelligence. The papers and demonstrations presented at the workshop covered a broad range of disciplines within AI including knowledge representation, machine learning, natural language processing, and online social media analytics. From an application perspective, presentations addressed topics in epidemiology, environmental and public health informatics, disease surveillance, health behavior monitoring, and disaster management.

One of the main themes of this workshop was the

exploration and monitoring of online social media (for example, Twitter) to analyze behavioral patterns. Models of behavior were used to enhance forecasting, guide decision making, enable situational awareness, and inform response strategies. The workshop also included three invited talks. Gregory Cooper (professor of biomedical informatics and of intelligent systems at the University of Pittsburgh) gave a presentation on detecting and characterizing outbreaks of infectious disease from clinical reports using Bayesian methods. Taha Kass-Hout (FDA chief health informatics officer, CTO, and director, Office of Informatics and Technology Innovation) described big data Initiatives at the FDA. Mark Dredze (assistant research professor in computer science at Johns Hopkins University and research scientist at the Human Language Technology Center of Excellence) presented uses of social media data for public health with an emphasis on the importance of online health-behavior monitoring and analysis.

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-15-15.

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

J. Christopher Beck is an associate professor in the Department of Mechanical and Industrial Engineering, University of Toronto.

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

Adi Botea is a research scientist at IBM Research, Dublin, Ireland.

Cornelia Caragea is an assistant professor in the Department of Computer Science and Engineering at University of North Texas.

Chi-hung Chi is from the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia.

Theodoros Damoulas is a research assistant professor in the Center for Urban Science and Progress at New York University.

Bistra Dilkina is an assistant professor at the College of Computing, Georgia Institute of Technology.

Eric Eaton is a faculty member at the University of Pennsylvania.

Pooyan Fazli is a postdoctoral fellow at Carnegie Mellon University.

Sam Ganzfried is a Ph.D. student in the Computer Science Department at Carnegie Mellon University.

Lee Giles is the David Reese Professor of Information Sciences and Technology at the Pennsylvania State University with appointments in the departments of Computer Science and Engineering, and Supply Chain and Information Systems.

Sébastien Guillet is a postdoctoral research fellow in computer science at the Université du Québec à Chicoutimi.

Robert Holte is a professor in the Computing Science Department at the University of Alberta.

Frank Hutter is an assistant professor at the University of Freiburg.

Thorsten Koch is a professor for software and algorithms for discrete optimization at TU Berlin and head of the Scientific Information Department at the Zuse Institute Berlin.

Matteo Leonetti is a postdoctoral fellow at the University of Texas at Austin.

Marius Lindauer is a research assistant at the University of Freiburg.

Marlos C. Machado is a Ph.D. student at the University of Alberta in Edmonton, Canada.

Yuri Malitsky is a researcher at IBM Research.

Gary Marcus is a professor of psychology and neuroscience at New York University.

Sebastian Meijer is a professor at KTH Royal Institute of Technology, Sweden, and Delft University of Technology, The Netherlands.

Francesca Rossi is a professor of computer science at the University of Padova, Italy.

Arash Shaban-Nejad, Ph.D., is a researcher at the School of Public Health, University of California, Berkeley.

Sylvie Thiebaut is a professor of computer science at the Australian National University and a research leader in the Optimisation Research Group at NICTA.

Manuela Veloso is the Herbert A. Simon University Professor of Computer Science at Carnegie Mellon University.

Toby Walsh is a professor at NICTA and University of New South Wales, Sydney, Australia.

Can Wang is from the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia.

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

Yu Zheng is from Microsoft Research, China.