

Reports on the 2016 IJCAI Workshop Series

Edited by Biplav Srivastava, Gita Sukthankar

■ *This report summarizes the 39 workshops held as part of the 2016 International Joint Conference on Artificial Intelligence, held in New York from July 9–11.*

The IJCAI 2016 Workshop Program was held in New York, New York, from July 9–11. It included 27 full-day and 12 half-day workshops covering a broad range of topics. Some workshops focused on recent improvements in computational frameworks such as statistical relational models and algorithmic game theory. The two deep learning workshops in particular generated a great deal of excitement with more than 200 attendees combined. However the majority of workshops centered on application areas including multiagent systems, natural language processing, cognitive computing, and social media analysis. The usage of artificial intelligence within the biological and health sciences was viewed as an important theme, and there were four workshops in that area. In addition to the specialized workshops, there was a workshop on ethics in artificial intelligence, a broad area of concern for most artificial intelligence researchers. Although most of the workshops were either completely new or had only been offered for a few years, one workshop (Qualitative Reasoning) celebrated its 29th year. Summaries of the workshops were presented at two evening sessions during the main conference for the benefit of attendees who could not attend the workshop days and are also posted online.¹

Bridging the Gap Between Human and Automated Reasoning

Human reasoning or the psychology of deduction is well researched in cognitive science. Automated deduction, on the other hand mainly focuses on the automated proof search in logical calculi. This workshop offered a platform for approaches coupling these two areas. The program covered the following themes: spatial cognition, theorem proving, and natural games; logic programming approaches to model human reasoning; syllogistic reasoning; computational models for human reasoning; benchmarks for commonsense reasoning; argumentation; automation of human inference.

During active exchange, it was remarked that even though the number of theories that claim to explain parts of human reasoning is strongly increasing, only few comparisons on common data sets do exist. Therefore, it was proposed that the field of human reasoning would greatly benefit from benchmarks and competitions similar to the ones in automated theorem proving. Benchmarks and a first competition are planned for the next Bridging the Gap Between Human and Automated Reasoning workshop.

The second Bridging the Gap Between Human and Automated Reasoning workshop was organized by Ulrich Furbach, Steffen Hölldobler, Marco Ragni, and Natarajan Shankar.

Computer Games Workshop

More than 30 attendees came together to interact about the advancements of AI in game-playing programs at the fifth Computer Games workshop. There were 11 presentations about different AI techniques and a variety of games. The games addressed were Breakthrough, Space Navigator, Same Game, One Night Ultimate Werewolf, General Game Playing, Atari games, Go, and Hex. The AI techniques included Monte Carlo tree search (MCTS), alpha-beta, nested rollout policy adaptation, deep learning, reinforcement learning, and natural language classifier. The most popular algorithms were MCTS and Deep Learning. The organizers plan to continue next year with a sixth Computer Games workshop.

The workshop was organized by Tristan Cazenave, Mark H. M. Winands, and Stefan Edelkamp.

Goal Reasoning

Goals are a unifying structure across the variety of intelligent systems, serving in the management of long-term behavior, anticipating the future, selecting among priorities, committing to action, generating expectations, assessing trade-offs, resolving the events, and learning from experience. The workshop on Goal Reasoning, the fourth in a series, was organized by Mark “Mak” Roberts (Naval Research Labora-

tory), Daniel Borrajo (Universidad Carlos III), Michael T. Cox (Wright State University), and Neil Yorke-Smith (American University of Beirut).

Participants discussed new formal models of goal reasoning, recognition, and representation; and explored the relationships with MDPs, automated planning, and HCI. Applications included a goal-reasoning algorithm operational in a \$100,000 UUV test. David Aha (NRL) and Sebastian Sardina (RMIT) gave invited talks.

Two of the major outcomes of the workshop were to challenge the assumption of static, user-provided goals, and to highlight the connection with BDI agent systems. After further revision and review, selected papers will appear in *AI Communications*.

Social Influence Analysis

The Second International Workshop on Social Influence Analysis (SocInf2016) featured the presentation of peer-reviewed accepted papers in which diverse social networks such as Twitter and Pinterest, hypergraphs, and even small groups (business meetings, group discussion) were used as case studies. Additionally, SocInf2016 included two invited keynotes: Big Network Analysis — Algorithms and Applications (by Jie Tang) and Negative Social Influence in Online Discussions (by Justin Cheng). These invited talks generated a great deal of audience interest. The workshop also hosted a contest sponsored by Alibaba Tianchi named “Brick-and-Mortar Store Recommendation with Budget Constraints” with \$10,000 in prizes.

The major outcomes of the workshop included the identification of different research gaps in the field and discussion of possible approaches to different aspects of social influence analysis.

The workshop was organized by Marcelo Armentano, Ariel Monteserin, Jie Tang, and Virginia Yannibelli.

Ethics for Artificial Intelligence

The Ethics for AI workshop featured the presentation of accepted peer-reviewed papers on a diverse set of topics concerning ethical and legal issues in artificial intelligence. Its aim was bringing together approaches from many disciplines and diverse points of view and fostering discussion to clarify questions and cooperative debate in finding ways forward. The disciplines of computer science, law, philosophy, economics, and social sciences were represented. Included were empirical approaches, for example, lab work furthering understanding of the ethical import of concepts such as “transparency”; as well as more theoretical work, such as work examining the relationship between ethical frameworks, law, and AI, how moral decisions are made, how best to think about central issues such as “autonomy.”

The workshop considered immediate concrete

issues in AI such as models for developing the regulation of autonomous vehicles, creating norms for the regulation of autonomous trading agents, as well as more speculative future projections of AI, the issue of superintelligence, its likelihood, and ethical import. The opinion was voiced that progress on these complex and important problems is being made.

The workshop was organized by Michael Wooldridge, Peter Millican, Christopher Megone, and Paula Boddington.

Computational Models of Natural Argument

In its 16th year, the Computational Models of Natural Argument workshop series² serves the argument and computation research community with a focus on natural argumentation, where naturalness may include expression in text, multimedia, or graphics; use of rhetorical devices; or taking into account characteristics of the audience such as affect. A distinctive feature of the workshop is its aim to reach different communities and perspectives. It has been colocated throughout the years with generalist AI conferences as well as more specialized meetings on multiagent systems, user modeling, and legal informatics.

The 2016 workshop featured a keynote talk and presentations of six papers and two research abstracts. Topics included models of argumentation for decision making, political analysis, and intelligence analysis; models of biomedical argumentation in research journals and popular media; annotation of rhetorical figures; and unique characteristics of argumentation in social media and in face-to-face dialogue.

The workshop was organized by Floris Bex, Floriana Grasso, and Nancy Green.

Interactive Machine Learning: Connecting Humans and Machines

The Interactive Machine Learning workshop focused on the design and analysis of algorithms that facilitate machine learning with the help of human interaction. During the workshop, 13 papers were presented covering a variety of topics such as human-machine teaching strategies, collaborative robot learning, topic models, and systems operations. There were four invited talks selected to span the breadth of research in this field, including presentations on reinforcement learning from users, active learning, cognitive science, and interactive robotics. A highlight of the workshop was a lively panel discussion on the realities of interactive machine learning, which focused on the real-world challenges of autonomous and interactive learning systems.

Some of the important talking points from the workshop included the importance of model trans-

parency, the influence of interaction modalities and the need to ask good questions for efficient learning. In future workshops, we hope to explore these ideas in greater detail as well as have better integration with the cognitive science and human-computer interaction community.

The workshop was organized by Kaushik Subramanian, Heni Ben Amor, Charles Isbell, and Andrea Thomaz.

Advances in Preference Handling

The 10th Multidisciplinary Workshop on Advances in Preference Handling, which brought together researchers from numerous subfields who are interested in computational aspects of preference handling, aims at discussing novel and emerging research on preferences, and provides an opportunity for cross-fertilization between fields. The program comprised 11 technical presentations and an invited talk by Vincent Conitzer on “Mechanism Design in Data-Rich Environments.” The papers are available on the workshop website.³

Preferences are a central concept of decision making and are used in fields such as AI, databases, and human-computer interaction. Preference-based systems do not return arbitrary responses, but choose results that respect users’ preferences. Embedding morality when handling preferences and dealing with the potential and risks of big data were identified as challenging endeavors for the future.

The workshop was organized by Markus Endres, Nicholas Mattei, and Andreas Pfandler.

Biomedical Informatics with Optimization and Machine Learning

The Biomedical Informatics with Optimization and Machine Learning workshop attracted a total of 25 submissions (full paper and short abstract tracks combined), that went through a peer-review process. Fifteen were accepted for oral or spotlight/poster presentations. Five keynote speakers from diverse backgrounds were invited, and more than 40 people attended the workshop. Best paper awards were chosen, with sponsorship by Microsoft Research. The workshop catalyzed synergies among biomedical informatics, machine learning, and optimization.

The workshop meets the compelling demand for novel machine learning, data mining, and optimization algorithms to specifically tackle the unique challenges associated with biomedical and health-care data. By reviewing and discussing the recent major breakthroughs in machine learning and optimization, it fosters idea exchanges among applied mathematicians, computer scientists, bioinformaticians,

computational biologists, industrial engineers, clinicians, and health-care researchers.

The workshop was organized by Zhangyang Wang, Yang Shen, Shuai Huang, and Jiayu Zhou.

Language Sense on Computers

The second Language Sense on Computers workshop was convened with the aim of gathering NLP specialists and linguists who work on the interesting phenomena and unsolved problems of language.

The workshop featured 13 presentations, which covered various topics including narratology and plot recognition, expressions for describing taste, word ordering tendencies, and approaches to metaphor processing. There were also papers oriented toward practical applications: one about recognizing Cockney rhyming for cyberbullying detection, one on automatic common sense ontology expansion, and one that introduced improved methods for affect recognition in dialogues. Two presentations were related to elderly-care solutions. Researchers from Japan presented work on analyzing a daily task (shopping) and considered the role of communication bots in therapy. Two papers addressed and tried to answer the following exploratory challenge problems: “Can computers write poetry?” and “Can computers predict the future?” The conclusion of the workshop was that difficult problems, although still being very hard to solve, show the weakest points of current natural language processing. We agreed that the challenges are definitely worth taking.

The workshop was organized by Akinori Abe and Rafal Rzepka.

AI for Synthetic Biology

Synthetic biology is the systematic design and engineering of biological systems; it holds the potential for revolutionary advances, for example, in medicine and environmental remediation. Unfortunately this potential is unreachable with the current practice of designing novel organisms at the DNA level, which is similar to programming in assembly code. Synthetic biology has reached a complexity barrier that AI can help to overcome, just like it did with programming languages that allow computer scientists to think and operate at higher abstraction levels while automating the translation to bits.

The AI for Synthetic Biology workshop aimed to cross-pollinate the two research fields. Thanks to the Bio-Design Automation Consortium and Raytheon BBN Technologies, the workshop provided funds to support synthetic biologists’ travel to ensure diversity of the audience. Invited talks introduced synthetic biology to AI researchers and highlighted areas where AI addresses synthetic biology challenges.

Contributed talks covered a diverse range of AI techniques applied to synthetic biology topics.

Synthetic biology is a rich domain for AI with many places for AI to make an impact. We hope this is the first of many workshops on this topic at AI venues that help develop collaborations between the two communities. The workshop was organized by Fusun Yaman, Aaron Adler, and June Medford.

Artificial Intelligence for Knowledge Management

The workshop covered various facets of knowledge management such as ontologies applied to access complex data in bioimaging or to local self-government, conceptual navigation and e-learning systems that use polyadic formal concepts, automatic selection of talents for collaborative new product design, self-organizing maps applied to RFM analysis in the packaging industry, selection of software for business intelligence, platforms for city dwellers, AI enhancing intelligent use of energy, segmentation of social networks, and semantic reasoning in knowledge management systems.

The invited talk was devoted to learning ontologies for medical data. Janusz Wojtusiak, associate professor of health informatics at George Mason University, Fairfax, recalled various machine-learning techniques and highlighted the importance of knowledge and context for deep exploring of medical data.

The proceedings of the workshop are available online.⁴ Selected, extended papers from this workshop will be published by Springer in the *Advances in Information and Communication Technology* series. The workshop was organized by Eunika Mercier-Laurent, Mieczyslaw Owoc, and Gulgun Kayakutlu.

Knowledge Discovery in Health-Care Data

The goal of the Knowledge Discovery in Health-Care workshop was to foster discussion and present progress on research efforts that leverage large amounts of observational data (clinical, biological, physiological) to expedite discovery in medicine. The workshop was intended to encourage a cross-disciplinary exchange of ideas between medical researchers and the artificial intelligence community. Artificial intelligence and machine-learning approaches hold the potential to reveal not readily apparent, hidden information in biological and medical health-care data sets. The results of such discoveries can aid the development of novel diagnostic and prognostic tests, inform descriptive, predictive, and prescriptive analytics, and guide hypothesis generation.

One of the highlights of the workshop was five world-renowned keynote speakers delivering four invited talks. The workshop also featured six paper

presentations and five posters that were accepted after the peer-review process. Participants were from Slovenia, United States, United Kingdom, and Japan.

The workshop was organized by Marzieh Nabi, Jonathan Rubin, Ali Shojaie, Ary L. Goldberger, Madalena Damásio da Costa, and Daniel G. Bobrow.

Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis

The International Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis⁵ was organized into two sections, one dedicated to auctions, mechanism design, and Walrasian equilibria; while the other was focused on problems related to power TAC. Following tradition, the workshop was associated with the Trading Agent Competition (TAC) and hosted the award ceremony for the two TAC 2016 tracks: Power and Ad Exchange TAC.

Mariano Schain, senior software engineer at Google, gave a fascinating talk on the TAC Ad Exchange Game. The workshop was organized by Sofia Ceppi, Chen Hajaj, Valentin Robu, Ioannis A. Vetsikas, and Esther David.

General Game Playing

AI researchers have for decades built game-playing agents capable of matching wits with the strongest humans in the world in board games like Chess and Go and video games such as StarCraft and Pac-Man. Research into general game playing takes this to the next level: to build intelligent software agents that can, given the rules of any game, automatically learn a strategy for playing that game at an expert level without human intervention.

Results discussed and achieved at the Fifth International General Game Playing workshop included boosting the efficiency of reasoning about game rules, general game playing with imperfect information, Monte Carlo tree search, and general video game playing. Selected papers from this workshop and the Computer Games workshop will be jointly published by Springer.

The workshop was organized by Stephan Schiffel, Julian Togelius, and Michael Thielscher.

Deep Learning for Artificial Intelligence

The Deep Learning for Artificial Intelligence workshop featured the presentation of peer-reviewed papers that focused on integrating deep learning (DL) with AI techniques (for example, for tasks involving the use of symbolic representations and related inference methods). Additionally, two invited speakers presented on DLAI-related topics: deep reinforce-

ment learning for robotics and explainable AI.

The papers focused on a variety of analysis (for example, classification) and synthesis tasks (for example, design). They also varied by application category (for example, natural language understanding, agent control, decision aids) and by their use of trained neural networks to assist with specific reasoning tasks (for example, solving Winograd schemas, goal selection, game move selection, object prediction, activity recognition, customer churn prediction). In summary, this workshop provided a forum for exploring how deep learning methods might interact and inform AI processes.

The workshop was organized by David W. Aha, Yianis Aloimonos, Andrew S. Gordon, and Alan Wagner.

Knowledge-Based Techniques for Problem Solving and Reasoning

The Knowledge-Based Techniques for Problem Solving and Reasoning workshop attempted to bridge the gap between knowledge representation communities (focusing on expressivity and semantics of models) and problem solving communities (focusing on efficient problem solving). It was inspired by area-specific workshops such as Knowledge Engineering for Planning and Scheduling and Constraint Modeling and Reformulation and the now discontinued Symposium on Abstraction, Reformulation, and Approximation.

The full-day workshop featured presentations of 10 peer-reviewed papers and an invited talk “The Modeling Beauty of Constraint Solving” given by Veronica Dahl. All papers are available through CEUR Workshop Proceedings. The topics of presentations varied from natural language processing over diagnosis and robotics up to search and planning. The papers showed the importance of problem modeling for efficient problem solving and confirmed the existing gap between the communities. Hence similar events in future seem desirable as a communication platform in particular at general AI conferences, where the knowledge representation and problem solving communities naturally meet.

The workshop was organized by Roman Barták, Lee McCluskey, and Enrico Pontelli.

Multiagent Path Finding

The Second International Workshop on Multiagent Path Finding focused on recent advances in multiagent path finding, particularly in the domain of optimal and near optimal methods.

Participants presented several new approaches to multiagent path finding, including using ad hoc group formation to ease reactive navigation for large groups of agents, any-angle planning for multiagent

systems, and structuring the environment to make planning easier. Invited speakers gave reviews of the state of the field.

The conclusion of a group discussion was that there are a large number of known algorithms, but our understanding of their benefits and drawbacks is insufficient. The participants committed to producing a good set of standard benchmarks to aid comparisons between approaches.

The workshop was organized by Howie Choset, Sven Koenig, and Glenn Wagner.

Computational Modeling of Attitudes and Sentiment Analysis

The combined workshops on Computational Modeling of Attitudes and Sentiment Analysis where AI Meets Psychology focused on the intersection between sentiment analysis and psychology and represented experts in both the psychological and cognitive sciences along with experts in computer science. Collectively, the papers illustrated the varied approaches to understanding sentiment and attitudes in real human data, to include theoretic models of human information processing and advanced machine-learning approaches for complex feature extraction of human-generated data. The types of human-generated data ranged from text to speech and audiovisual data. The keynote speakers represented world-renowned researchers from machine learning (Björn W. Schuller) and social psychology (Russell Fazio). The final round-table discussion brought novel insights to both the psychological and the computer science participants and a sense that there was room for further cross-fertilization to understand human behavior and to provide more sophisticated machine-learning applications.

The two workshops were organized by Sivaji Bandyopadhyay, Erik Cambria, Dipankar Das, Braja Gopal Patra, Kiran Lakkaraju, Mark Orr, and Samarth Swarup.

Semantic Machine Learning

The 2016 Semantic Machine Learning workshop focused on the latest advances and research challenges in fusing data and functional and domain semantics with traditional machine learning toward enhancing its performance. The workshop featured two invited keynotes, four paper presentations, and a panel discussion that were well received. The keynotes highlighted the importance of unsupervised learning and illustrated diverse ways to formalize and incorporate domain semantics into it, while the panel on challenges and potential directions to improve machine learning with semantics facilitated diverse perspectives on semantic machine learning.

The new results from papers demonstrated

approaches ranging from incorporating structured KB's into machine learning to exploiting deep learning for extracting domain semantics. We identified key research priorities for the future as: knowledge representation with evolution, and machine learning with explanation. The workshop proceedings, presentation slides, and video recordings are available online.⁶

The workshop was organized by Rajaraman Kanagasabai, Ahsan Morshed, and Hemant Purohit.

Natural Language Processing Meets Journalism

Natural language processing has matured over the years to the point where a suite of technologies is available to cope with many problems raised by the contemporary need of global information. Consequently, it is now time for NLP to become engaged as an active partner for both journalists and readers.

The Natural Language Processing Meets Journalism workshop attracted the interest of researchers both in computational linguistics and journalism and of professionals in the news production system. The primary goal was to have a forum in which it would be possible to share and discuss advancements in natural language processing and real needs in the field of journalism. At this workshop there were papers that presented technologies that go deep into the substance of a piece of news, as a savvy journalist or an alert reader will do.

The workshop was organized by Octavian Popescu, Carlo Strapparava, and Larry Birnbaum.

Advances in Bioinformatics and Artificial Intelligence: Bridging the Gap

This one-day workshop brought together scholars and practitioners active in AI-driven bioinformatics, to present and discuss their research, share their knowledge and experiences, and discuss the current state of the art and the future improvements to advance the intelligent practice of computational biology.

The 2016 program comprised one keynote speaker, two invited speakers and seven paper presentations. The details are accessible online.⁷ The presentation topics were divided into two areas: biology-inspired computation and computation providing new insight into biology. The workshop covered the broad scope of AI and bioinformatics, including machine learning, knowledge representation, natural language processing, web mining, comparative genomics, system biology, and networks.

There is a plan for for a future workshop to discuss personalized medicine. The workshop was organized by Abdoulaye Baniré Diallo, Engelbert Mephu Guifo, and Mohamed Zaki.

Algorithmic Game Theory

The workshop on Algorithmic Game Theory featured the presentation of 11 technical papers in various topics of algorithmic game theory (mainly networks and social choice). In addition, there were an hour-long invited talk by David Parkes and a short rump session for advertising AGT-related papers in the main IJCAI session. The workshop brought together AI researchers that are interested in strategy and game theory, but from different perspectives: using computational techniques to games, efficiently computing and implementing solution concepts, and exploring alternative notions of individual and social utility.

The workshop was organized by Georgios Chalkiadakis, Nicola Gatti, Reshef Meir, and Carmine Ventre.

Statistical Relational AI

The purpose of the Sixth International Workshop on Statistical Relational AI workshop series is to bring together researchers and practitioners from two fields: logical (or relational) AI and probabilistic (or statistical) AI. Until recently, research in them has progressed independently with little or no interaction. The workshops instead provide a big picture view on AI. It is the study and design of intelligent agents that act in noisy worlds composed of objects and relations among the objects.

The 2016 workshop featured three invited talks, by William Cohen, on “TensorLog: A Differentiable Deductive Database,” by Daniel Lowd, on “Adversarial Statistical Relational AI,” and by Percy Liang, on “Querying Unnormalized and Incomplete Knowledge Bases.” The workshop had 25 accepted papers, which were presented as spotlight talks and posters. The workshop details are available online.⁸

The workshop was organized by Guy Van den Broeck, Mathias Niepert, Sebastian Riedel, and David Poole.

Deep Reinforcement Learning: Frontiers and Challenges

Deep learning and reinforcement learning are among the most promising machine-learning methods today. This workshop focused on applications of deep learning for representation learning in reinforcement learning and applications of reinforcement learning for more efficient deep learning. The primary motivation of organizing the workshop here was to involve the IJCAI community in this research drive.

The workshop featured seven keynote talks by leading researchers in the field: Pieter Abbeel, Remi Munos, Joelle Pineau, Doina Precup, David Silver, Satinder Singh, and Peter Stone. Speakers covered various topics including deep RL for games, NLP, robotics, and RL for deep learning. The workshop also featured 10 contributed papers and a panel discussion on

research challenges in deep RL. There was significant interest from the audience and several expressed the hope that this will continue next year as well.

The workshop was organized by Sarath Chandar, Sridhar Mahadevan, Balaraman Ravindran, and Gerald Tesaro.

Natural Language Processing for Social Media

The goal of the research discussed at the Fourth International Workshop on Natural Language Processing for Social Media was to enhance social computing with AI and natural language processing (NLP). Presentations focused on NLP problems using information extracted or learned from social networks and social media. New research problems related to both social computing and NLP were highlighted during the workshop.

The workshop was organized into four parts. First, Yuheng Hu (University of Illinois at Chicago) delivered an excellent keynote speech on event analysis in social media. His talk received great feedback and brought lively discussions among the participants on the insights of people’s engagement with events and the tweeting behaviors during engaged events. Second, sentiment analysis using AI, especially machine-learning techniques, was one of the key workshop topics. Third, deep learning was mentioned by every presentation. Fourth, due to the importance of benchmark datasets, SocialNLP encouraged data papers to share resource and data creation and preliminary analysis. Two interesting data track papers were accepted this year, one on Hindi-English Mixing, and another on Moroccan Arabic code switching. The organizers have maintained a modest size with 6 full paper presentations and a total of 20 to 25 participants.

The workshop was organized by Jane Yung-jen Hsu, Lun-Wei Ku, and Cheng-Te Li.

Qualitative Reasoning

Qualitative Reasoning workshops discuss research on automated understanding of the world derived from incomplete, imprecise, or uncertain data, realized as cognitive systems capable of knowledge-level interaction (with humans in the loop).

This year’s workshop hosted an inspiring invited talk by Diedrich Wolter on qualitative spatial reasoning, followed by 14 presentations of ongoing research. The presentations focused on topics including conceptual modeling and simulation for education (learning); diagnosis and decision making (for example, environmental problems); explanatory models for health, biodegradation, and science; order-of-magnitude reasoning (for business and marketing); human and (physical) robot interaction during gaming; and, of course, qualitative spatial rea-

soning. All contributions can be found on the website of the workshop.⁹

A key observation throughout the workshop was that contemporary challenges typically concern multidimensional problems, which require semantic interoperability of miscellaneous representations and algorithms.

The workshop was organized by Bert Bredeweg, Kamal Kansou, and Matthew Klenk.

Cognitive Knowledge Acquisition and Applications

Motivated by the goal of developing cognitive systems that learn, reason, and interact naturally with humans, the Cognitive Knowledge Acquisition and Applications workshop series¹⁰ offers an informal setting that promotes lively discussion on work that bridges cognitive psychology and artificial intelligence. Emphasis is placed on developing knowledge representations and acquisition processes that allow a cognitive system to explain the inferences that it draws, while also being able to accommodate user feedback to adapt and improve its workings.

The workshop featured a keynote talk by Ernest Davis on building a corpus of texts tagged with commonsense inferences and contributed talks on learning knowledge for commonsense psychology, on robots learning to ground text on their visual inputs, on recommending films while explaining why, on using Minsky's critic networks for common sense, on enhancing expert-developed taxonomies, and on achieving more flexibility than static dictionaries through learning.

The workshop was organized by Loizos Michael and Erik T. Mueller.

Interactions with Mixed Agent Types

The Interactions with Mixed Agent Types workshop centered on fostering a reach discussion about how to design intelligent agents capable of interacting with other agents of different types and whose objectives, learning dynamics, and representation of the world are

unknown. It featured the presentation of peer-reviewed papers that focused on different sorts of mixed agent type interactions. It was composed of two sections, one dedicated to modeling the strategic reasoning of the opponents under different circumstances and the other dedicated to planning and optimization in mixed agent type settings.

In conclusion, a growing number of investigations consider a variety of interacting agents. However, much effort is still needed to close the gap between the state of the art and heterogeneous multiagent systems. With this in mind, the community still needs to assemble diverse perspectives to promote a robust understanding of agent mix.

The workshop was organized by Enrique Munoz de Cote, Long Tran-Thanh, Christopher Amato, and Prashant Doshi.

Ontologies and Logic Programming for Query Answering

The Ontologies and Logic Programming for Query Answering workshop featured the presentation of accepted papers that underwent a peer-review process. Some papers addressed query answering while taking ontologies into account with new results on temporal OBDA and on a unified framework for inconsistency-tolerant query answering; other papers were devoted to non-monotonic reasoning, inconsistency handling and expressing default negation in ontologies with new results on mapping data to ontologies using existential ASP, on qualitative disjunctive logic programs for ontology mapping, and on existential ASP for computing repairs. Additionally an invited speaker focused on query rewriting for ontology-based query answering.

The workshop gave rise to rich discussions and an open issue emerged: how to define fragments of existential ASP covering lightweight ontological languages while keeping decidability and efficiency?

The workshop was organized by Odile Papini, Salem Benferhat, Laurent Garcia, and Marie-Laure Mugnier.

Closing the Cognitive Loop for Human-Aware AI

The Third Closing the Cognitive Loop for Human-Aware AI workshop featured the presentation of accepted papers that went through a peer-review process and focused on successful addressing of human-in-the-loop issues. Additionally, two invited speakers presented deployed applications of AI that interact closely with humans — cognitive assistance for data scientists, and intelligent control of crowdsourcing applications.

One of the major outcomes of the workshop was the realization that each AI system that was presented in the workshop featured a unique set of interaction challenges that needed to be overcome in order to make those systems human aware. This sparked a discussion on the major problem pillars for human-aware AI, which include topics such as explanation of decisions; interpretability of the decision process; efficient and time-sensitive context transfer; division of skills and labor; and legal and ethical issues. The next iteration of the workshop will focus more specifically on these problem pillars for human-aware AI.

The workshop was organized by Karthik Talamadupula, Shirin Sohrabi, and Murray Campbell.

Notes

1. bit.ly/2atCpD9.
2. cmna.info.
3. www.mpref-2016.preflib.org.
4. ifipgroup.com/AI4KMProceedings2016.pdf.
5. www.sofiaceppi.com/AMECTADA2016.
6. datam.i2r.a-star.edu.sg/sml16.
7. bioinfo.uqam.ca/IJCAI_BAI2016.
8. www.starai.org.
9. ivi.fnwi.uva.nl/tcs/QRgroup/qr16/index.html.
10. cognitum.ws.

Biplav Srivastava is a research staff member at IBM Thomas J. Watson Research Center, USA.

Gita Sukthankar is an associate professor at University of Central Florida in Orlando.