

Reports of the Workshops of the 32nd AAAI Conference on Artificial Intelligence

Bruno Bouchard, Kevin Bouchard, Noam Brown, Niyati Chhaya, Eitan Farchi, Sebastien Gaboury, Christopher Geib, Amelie Gyrard, Kokil Jaidka, Sarah Keren, Roni Khardon, Parisa Kordjamshidi, David Martinez, Nicholas Mattei, Martin Michalowski, Reuth Mirsky, Joseph Osborn, Cem Sahin, Arash Shaban-Nejad, Onn Shehory, Amit Sheth, Ilan Shimshoni, Howie Shrobe, Arunesh Sinha, Atanu R. Sinha, Biplav Srivastava, William Streilein, Georgios Theodorou, K. Brent Venable, Neal Wagner, Anna Zamansky

■ *The AAAI-18 workshop program offered 15 workshops covering a wide range of topics in AI. The workshops were held February 2–3, 2018, at the Hilton New Orleans Riverside in New Orleans, Louisiana. This report contains summaries of the Affective Content Analysis workshop; the Artificial Intelligence Applied to Assistive Technologies and Smart Environments; the AI and Marketing Science workshop; the Artificial Intelligence for Cyber Security workshop; the AI for Imperfect-Information Games; the Declarative Learning Based Programming workshop; the Engineering Dependable and Secure Machine Learning Systems workshop; the Health Intelligence workshop; the Knowledge Extraction from Games workshop; the Plan, Activity, and Intent Recognition workshop; the Planning and Inference workshop; the Preference Handling workshop; the Reasoning and Learning for Human-Machine Dialogues workshop; and the the AI Enhanced Internet of Things Data Processing for Intelligent Applications workshop.*

Affective Content Analysis

The Affective Content Analysis workshop was conducted as an interdisciplinary platform to stimulate cross-disciplinary discussions on affect in content and to more deeply involve the AI and ML community in the open problems in affective content analysis, with a special focus on affect in language and text. Affective content analysis in this context refers to the interdisciplinary research space of computational linguistics, psycholinguistics, consumer psychology, and human-computer interaction (HCI) with respect to the various forms of online communication. The number of workshops and conferences related to affective computing has been growing, which points to the importance of the research problem, as well as the timeliness of this workshop for the AI community.

Affective computing has traditionally focused on modeling human reactions using multimodal sensor data, but not text.

Sentiment and emotion analysis, on the other hand, has been applied to both text and multimodal datasets, but this research has been limited to quantifying well-defined human reactions. Affect analysis (that is, techniques and applications to understand the experience of an emotion) in the context of language and text is an up-and-coming research space. Work on affect analysis in language and text spans many research communities: computational linguistics, consumer psychology, HCI, marketing science, and cognitive science. Computational linguists study how language evokes, as well as expresses, emotion. Consumer psychology examines human affect by drawing upon grounded psychological theories of human behavior. The HCI community studies human responses as a part of user experience evaluation. Computational models for consumer psychology theories present a huge opportunity to guide the construction of intelligent systems that understand human reactions, and tools from linguistics and machine learning can provide attractive methods to fulfill those opportunities. Models of affect have recently been adapted for social media platforms, enabling new approaches to understanding user's opinions, intentions, and expressions.

The workshop focused on the analysis of emotions, sentiments, and attitudes in textual, visual, and multimodal content for applications in psychology, consumer behavior, language understanding, and computer vision. Besides original research presentations and posters, the workshop also hosted a range of keynote speakers who highlighted the state of the art in affective computing in a range of fields.

James Pennebaker from the University of Texas Austin provided evidence from a series of studies about how affect and emotion can be mined from the words used by people in everyday life. Dipankar Chakravarti from Virginia Tech discussed some of the challenges involved in affective analysis of text for consumer behavior, especially noting the differences between the experience and the expression of affect. Bjoern Schuller from the University of Augsburg, Germany, provided insight into a range of applications of affect analysis from speech, music, and audio. Rajesh Bagchi from Virginia Tech shared work in the space of consumer psychology and marketing science, focusing on the affective processing of information and its relationship to consumer behavior. Cristian Danescu-Niculescu-Mizil talked about the application of affective computing in conversational dynamics, in group discussions, and with respect to the outcomes of decision-making discussions. Jennifer Healey from Intel discussed her work in multimodal affect analysis as a part of the cutting-edge research on building emotionally aware robots that can intuit and respond to human emotions.

The workshop ended with a panel discussion among the keynote speakers, moderated by the organizers, on the potential grounds for interdisci-

plinary collaborations, as well as venues for such events in future.

Niyati Chhaya, Kokil Jaidka, Lyle Unger, and P. Aanandan cochaired the workshop. This report was prepared by Niyati Chhaya and Kokil Jaidka. The workshop papers were published in the AAAI digital library.

Artificial Intelligence and Marketing Science

The growth in online data across marketing, campaign, display, programmatic advertisements, and social platforms has focused the attention of AI and machine learning researchers on developing new and more efficient computational models and techniques. Among marketing science (MS) researchers, the emphasis has been on exploiting ML methods to address business problems in marketing resource optimization, managerial decision-making, competitive behavior modeling, deconstruction of consumer behavior, and campaign automation and optimization. These two vibrant research communities both investigate problems central to marketing, but publish in separate journals and conferences. The motivation for the workshop was to start the dialogue for bridging that divide. The desired, longer term outcome is the two communities benefiting from each other's research, problems, and insights, leading to the higher effectiveness of models and methods in both theory and application.

The workshop showcased four invited keynote speakers to explore a number of key themes: Craig Boutilier (Google Inc.), Vince Conitzer (Duke University), Harikesh Nair (Stanford University and JD.com), and Dave Weinstein (Adobe Systems Inc.). The 11 paper presentations included four posters.

Boutilier presented research focusing on sequential decision-making, a major area straddling both AI and MS. His presentation highlighted advances in Markov decision processes (MDPs) to tackle problems in online advertising related to understanding the long-term impact of advertisements. His recent work includes logistic MDP and stochastic action set MDP. The latter is useful when new, exogenous actions are to be considered (for example, a new advertising campaign). In the context of marketing decision-making, algorithms must address the double issue of partial observability and multiple goals, which is where MDPs and reinforcement learning become particularly relevant. Voice of the customer data, including information on user goals, can be mined to help with the goal-identification tasks important for MDPs.

Conitzer presented research that addresses the problem of pacing in selecting bids on a demand-side platform. Consider that a marketer asks an agent (the platform can also be the agent) to bid on its behalf subject to a daily budget. An important problem is



2 - 7 February

New Orleans, Louisiana, USA

Program Chairs

Sheila McIlraith (University of Toronto)

Kilian Weinberger (Cornell University)

www.aaai.org/aaai18



how to bid over the day, given a target specified by the marketer and a budget. The problem involves selecting the bid per impression, identifying the total number of bids, and spreading the bids over the day. The latter is pacing. Conitzer has established a pacing equilibrium bidding strategy. The work is useful for conducting auctions for several marketing actions and is important in focusing attention on the problem of resource allocation in marketing. Given the problem of endogeneity, another relevant task is finding exogenous variation in data. One presentation presented a method for doing so.

Nair discussed the approach that marketing science takes to research, with the goal of informing decision-making for business problems. At the start, a model of consumer behavior is built, informed both by economics and psychology. Aggregated across consumers a demand function is obtained, which has as arguments marketing actions. The actions are optimized for appropriate business objectives. The approach is driven by theory, drawing from received wisdom in consumer psychology and economics. The need for prediction and scalability is addressed by AI and ML, while the need to bring in theory-driven models is met by MS. A common platform requires building models from the ground up, rather than setting models up to learn from data. Recognizing the process by which data is generated becomes important for answering questions of causation. Rounding out the theme of understanding consumer behavior for modeling were methods for extracting brand perceptions from online data.

Weinstein presented the industry perspective on AI's future in addressing real-life marketing problems, including the broad area of advertising and promotions. Weinstein's deep experience across industries, his insider perspective, and his insights regarding the distinctions between first-party, second-party, and third-party data, and their different uses when combined, provide new grounds for research for both AI and MS communities. With new research opportunities arises the need for new thinking about methods and models, much of which can benefit from cross-pollination across AI and MS. Paper presentations provided an understanding of consumer behavior and response to adverts, useful input for marketers in making advertising decisions.

The speakers, presenters, and participants agreed that this workshop was a valuable experience in bringing together academic and industry researchers from both AI and MS. Attendees agreed that more workshops should be organized to build on this event, because the confluence of research is critical for AI and MS to advance both research and practice in this digital economy.

Hung Bui, Pradeep Chintagunta, S. Muthukrishnan, Tuomas Sandholm, Atanu R Sinha, and Georgios Theodorou served as cochairs of this workshop. This report was written by Atanu Sinha and

Georgios Theodorou. The workshop papers were published in the AAAI digital library.

Artificial Intelligence Applied to Assistive Technologies and Smart Environments

Recent progress in AI is reshaping the way we conceive of the world. With the proliferation of sensors and the lowering cost of smart devices, new subfields of AI have emerged, such as assistive technology and smart environments. A smart environment is a physical space, perhaps a living facility, equipped with sensors, actuators, and AI capabilities, which together provide smart services to the occupant. That kind of environment has the potential to enhance quality of life by providing assistance in the activities of daily living, rendering the environment itself a form of assistive technology. Environments such as these are particularly interesting for the support they can provide elders and impaired persons, improving their autonomy and reducing the need for caregivers. The market for assistive technologies is rapidly growing, reaching \$60.5 billion in 2018 in the US alone, with an expected growth of almost 6 percent annually in the next decade. Given this expansion, it has become a major sector in research and development.

Despite the growing interest in these technologies, however, they have not yet been widely adopted. Indeed, the impairments and particularities of users are so diverse that implementing solutions for well-being represents one of the major challenges of universal design. The goal of this workshop was to investigate new solutions to scientific problems in various topics related to artificial intelligence as applied in the domain of assistive technology and smart environments for persons with special needs.

The 2018 AAAI workshop brought together academic and industrial researchers from several subfields of AI. One main theme of the papers presented at the workshop was learning methods used to personalize the assistance to a specific user. In fact, this theme represents one major challenge in providing good assistance: assistive systems need to learn precisely the profile of each of its users in order to be able to help them effectively. Several papers presented learning approaches with validations and experiments on real data.

Another major theme was the development and testing of new assistive devices. The papers on this theme can be divided into two categories. The first concerned new assistive devices to help populations with physical impairments, such as smart arms for wheelchairs. The second focused on devices for persons with cognitive deficits.

The workshop participants discussed the ways in which assistive technologies can help create a better future for populations with special needs, touching too on how these technologies can benefit from

innovation in domains such as the Internet of Things. Participants shared the objective of developing methods to reduce the costs and increase the efficiency of assistive technology, and agreed that they would like to attend a future edition of the workshop.

Bruno Bouchard, Sebastien Gaboury, and Kevin Bouchard cochaired this workshop and prepared this report. The workshop papers were published in the AAAI digital library.

AI-Enhanced Internet of Things Data Processing for Intelligent Applications

The purpose of this workshop was to discuss how AI techniques can help consume data from IoT to build intelligent applications. The workshop brought together 35 academic researchers and industry practitioners from around the world interested in taking a multidisciplinary approach to advance the state of the art of IoT to solve real-world business and societal problems. The technical program consisted of eight peer-reviewed papers (four full presentations and four posters, along with lightning talks), two invited talks, and one panel.

The program began with an invited talk titled “AI for Social Good and IoT: With Applications to Wildlife Conservation,” by Milind Tambe (University of Southern California). Tambe explained his group’s work applying game theory to various security-related applications, including strategies for forest rangers for preventing poaching. In the specific problem discussed, forest rangers can employ drones armed with video cameras to detect and deter poachers. The solutions work to predict the location of traps, to coordinate patrollers and drones, and to use sensor monitoring and interdiction strategies for maximum effectiveness. The post-lunch session included an invited talk by Nimit Desai (IBM Research), who spoke about edge computing and its applications to health and safety for large numbers of the population. The talk described two applications of edge computing: Mesh Network Alerts (MNA), which enables citizens to receive severe weather information even when there is no connectivity; and Nightingale, which enables patients to receive personalized health predictions. MNA is a delay-tolerant peer-to-peer network of commonly available Android and iOS devices that is presently in use by more than a million users in emerging countries. The network is used to broadcast severe weather information in a trusted manner without requiring any networking infrastructure. Nightingale incorporates early ideas on a combination of distributed learning and transfer learning on mobile devices to provide personalized risk assessments to patients suffering from asthma, migraines, and other disorders. Although the risk assessments are based on weather,

user history, and user behavior data, the data never leaves the users’ devices so as to protect user privacy.

The paper presentations covered topics related to reasoning for traffic signal coordination, learning graphical models from IoT data and time series analysis, and analyzing images and audio for object detection. The concluding event of the day was a panel discussion led by Amit Sheth and Nimit Desai on focus areas for AI to make most of sensor data and application potential. The discussion involved the active participation of the audience. Speakers discussed the role and importance of semantics (semantic sensor network ontology and its relationship with domain-specific sensor metadata), machine learning at the edge, and challenges like adversarial ML and the need for abstraction to scale reasoning.

The event offered an invigorating day of technical exchange, both formal and informal. The attendees felt the need to build further research momentum on the topic at future AI conferences.

Payam Barnaghi, Amelie Gyrard, Amit Sheth, and Biplav Srivastava served as cochairs of the workshop and coauthored this report. The workshop papers were published in the AAAI digital library.

Artificial Intelligence for Cyber Security

The Artificial Intelligence for Cyber Security workshop focused on research and applications of AI to operational problems in cyber security, including machine learning, game theory, threat modeling, and automated and assistive reasoning. Talks focused on the application of AI to ensure the resilience of the Internet of Things (IoT) and mobile devices, with emphasis on those that involve human-machine interactions.

The workshop began with a keynote speech by Salvatore Stolfo (Columbia University), “An Overview of Machine Learning Applied to Security.” Stolfo began with a historical perspective on the early applications of AI to fraud detection and intrusion detection systems (IDS). He emphasized areas where machine learning can have an impact, such as in the role of generating cyber decoys and defending against attacks that exploit embedded device (IoT) vulnerabilities. He pointed out other areas where the field of AI can also have significant impact as it applies to anomaly detection and social network analysis relevant to cyber security.

The initial session featured talks on adversarial learning. The first paper presented an approach to generating adversarial malware examples by leveraging an approximation to the defender’s hidden classifier, a recurrent neural network (RNN). Results showed an enhanced ability to evade detection. The second paper presented a novel approach to combating adversarial deceit that leveraged a probabilistic Stackelberg game. Simulated results demonstrated

the method's effectiveness within the context of defender security patrols. A third paper explored the effectiveness of an enhanced support vector machine kernel against evasion attacks generated via gradient-descent (GD) methods. The authors also introduced a new method for improving attack success rate based on an improved GD method.

The next session focused on Internet of Things security and associated attacks. The first paper described a novel way to exploit IoT side channels to send messages that are not detected by existing IoT protocols. Additionally, the authors showed that the channel capacity for such messages can be mathematically bounded. The second paper focused on mission-aware network utilization. The authors proposed a system called RADMAX (risk and deadline-aware planning for maximum utility) that employs constraint programming to maximize overall network utilization such that mission-specific routing flows do not violate their deadlines or loss requirements.

The workshop continued with a panel discussion on the threats from emerging technologies, and included panelists Brett Meyers (CrowdStrike), Howard Shrobe (MIT), Trung Tran (University of Maryland), and Salvatore Stolfo (Columbia University). Panelists discussed the important security issues associated with ubiquitous insecure IoT devices. Particular emphasis was placed on how attackers can leverage these devices for mounting larger attacks against critical infrastructure and other targets. The panel then discussed the role that market forces and government regulations play in addressing these important issues, referencing Europe's general data protection regulation (GDPR).

The afternoon keynote from Trung Tran (University of Maryland) on approaches to cyber security stressed that because of the cost imbalance inherent in the attacker and defender security model, the advantage is with the attacker, who will never tire and who will eventually succeed in his or her attacks. Rather than focus on individual attacks, however, we should attempt to stop the attacker altogether. By letting the attacker in, in a controlled way, we can observe and learn their motivations and methods to better inform our defenses and also to preemptively act. Trung concluded with a call to action and encouraged the community to develop novel approaches that leverage game-theoretic and machine learning approaches to support automated responses.

The workshop issued its second challenge problem in as many years. This year's challenge invited participants to create novel machine learning algorithms to detect unreliable news articles. The challenge included a corpus of data containing real internet news articles labeled as reliable or unreliable according to a curated, open source listing of unreliable news sources. The winning paper compared the

performance of multiple natural language processing (NLP) algorithms toward this task. Results showed that hierarchical attention networks (HANs) are the most effective of the NLP algorithms examined and advocated for their use.

The final session featured research on data analysis methods. The first paper described an approach to audit analysis that achieved promising anomaly detection results by leveraging recurrent neural network (RNN) language models. A small amount of training data was needed to achieve complete recall. The second paper presented a multiview data analysis approach to attack identification that minimized false alarms by incorporating analyst feedback. The next paper described an approach to malware detection that leveraged a neural network trained on the byte sequence of the entire malicious executable. Results were improved over n-gram methods and those that focus on the program execution (PE) header. The final paper of the session presented an integrated approach to computer security data labeling that leveraged active learning and real-world user experience. Results from user studies were very favorable.

This was the third annual AI for Cyber Security workshop. The meeting had more than 50 attendees. It was cochaired by William Streilein, David Martinez, Cem Sahin, Howie Shrobe, Arunesh Sinha, and Neal Wagner, all of whom prepared this report. The workshop papers were published in the AAAI digital library.

Artificial Intelligence for Imperfect-Information Games

Imperfect-information games are used to model various strategic interactions involving hidden information such as negotiations, auctions, and security interactions (both physical and virtual). Because of the presence of hidden information, solving these games requires methods quite different from traditional games of perfect information like chess or Go. Although there has been considerable recent progress in a number of communities studying imperfect-information games, the techniques used by each community have remained relatively isolated despite their generality. There is ample opportunity for cross-pollination between these communities, leading to new applications of methods already popular in one community, or the creation of new techniques by building upon the methods already established in separate communities.

This workshop brought together researchers studying aspects of imperfect-information games such as two-player zero-sum games, security games, financial markets, and sequential social dilemmas. Researchers in these areas typically work in isolation from one another, despite the similarities both of the domains and of the techniques used to address them. The primary goal of the workshop was therefore to expose

researchers to advances being made in other subareas and to encourage collaboration between these typically isolated communities.

At the conclusion of the workshop, participants discussed future benchmark challenge problems for the field. The classic challenge problem had been poker, but in 2017 AI exceeded top human performance in this game. While no consensus was reached on a particular domain, there was agreement that the field should look beyond zero-sum games and ideally beyond recreational games toward more real-world applications.

Noam Brown, Marc Lanctot, and Haifeng Xu served as organizers of the workshop. This report was prepared by Noam Brown. The workshop papers were published in the AAAI digital library.

Declarative Learning-Based Programming

Currently, to solve the real-world problems in many areas such as cognitive sciences, biology, finance, and the social sciences, scientists think about data-driven solutions. However, current technologies offer cumbersome solutions along multiple dimensions. Such dimensions include interaction with raw data, the need for extensive programming, the need to exploit various learning techniques, and extensive experimental exploration for appropriate models. Declarative learning-based programming investigates how one can facilitate and simplify the design and development of intelligent real-world applications that engage in learning from data and reasoning based on knowledge. This form of programming highlights the challenges in making machine learning accessible both to various domain experts and to application programmers. Conventional programming languages have not been primarily designed to offer help for these challenges.

To achieve the goals of declarative learning-based programming, we need to go beyond classic machine learning and AI tools. We need innovative abstractions that enrich the existing solutions and frameworks with capabilities in specifying the requirements of the application at a high level of abstraction. We need abstractions that exploit expert knowledge in learning, that deal with uncertainty in data and knowledge, and that use representations to support flexible relational feature engineering, flexible reasoning, and structure learning. We need abstractions that encourage reuse, that combine and chain models, and that perform flexible inference on complex models or pipelines of decision-making. We need abstractions that integrate a range of learning and inference algorithms, that close the loop on moving from data to knowledge, that exploit knowledge to generate data, and that offer a unified programming environment in which to design application programs.



Over the years, the research community has attempted to address these problems from multiple perspectives, most notably various approaches based on probabilistic programming, logical programming, constrained conditional models, and other integrated paradigms such as probabilistic logical programming and statistical relational learning. In recent years, several deep learning tools have created easy-to-use abstractions for programming model configurations for deep architectures. We aim at motivating the need for further research toward a unified framework in this area based on these key existing paradigms. To this end, we investigated the required type of languages, representations, and computational models to support a declarative, learning-based programming paradigm.

The workshop brought together researchers from various areas related to relational databases, relational learning, probabilistic logical databases, probabilistic programming languages, and neural program interpreters. Parisa Kordjamshidi provided a brief overview of the aims and the challenges of the declarative, learning-based programming paradigm. Avi Pfeffer (Charles River Analytics) provided an invited talk on a very new deep probabilistic language called Scruff. Sebastian Riedel (University College London) discussed the recent research progress on neural program interpreters and applications to machine reading and reasoning. William Cohen (Carnegie Mellon University) discussed recent progress made on probabilistic logics and declarative statistical learning platforms. Eli Bingham from Uber's development team described the progress they have made on Pyro, a probabilistic programming language based on Python and PyTorch. Alex Ranter

(Stanford University) discussed recent results on Snorkel, a platform for rapid training and data creation with weak supervision. Two accepted papers were presented in the workshop, one by Golnoosh Farnadi on fairness-aware relational learning and another by Michelangelo Diligenti on the application of combining deep learning with prior knowledge for image classification. We held a panel discussion that included both invited speakers and organizing committee members. The discussion ranged over topics such as important directions for the future, the issues of declarative versus procedural programming paradigms, probabilistic logical versus deep learning programming platforms, and the missing components in the current research toward developing systems AI.

Parisa Kordjamshidi, Dan Roth, and Kristian Kersting served as cochairs of this workshop. This report was prepared by Parisa Kordjamshidi. The two accepted papers were published in the AAAI digital library.

Engineering Dependable and Secure Machine Learning Systems

Businesses and the society at large increasingly rely on machine learning solutions. Similar to other software systems, machine learning systems must meet certain requirements. Indeed, different system types may introduce different dependability and quality requirements. However, meeting reliability, quality, and security requirements in the context of ML requires new methodologies and new tools.

Standard notions of software quality and reliability such as deterministic functional correctness, black-box testing, code coverage, or traditional software debugging become practically irrelevant for ML systems. This irrelevancy follows from the non-deterministic nature of ML, from the reuse of high-quality implementations of ML algorithms, and from the lack of visibility into the semantics of learned models, for example, when deep learning methods are applied. This difference from traditional programming calls for novel methods to address quality and reliability challenges of ML systems.

Broad deployment of ML software in networked systems inevitably exposes that ML software to attacks. While classical security vulnerabilities are relevant, ML techniques have additional weaknesses, some already known (for example, sensitivity to training data manipulation), and some yet to be discovered. Hence, there is a need for research as well as practical solutions to ML security problems.

This workshop focused on such topics. It included original contributions addressing problems and solutions related to dependability, to quality assurance, and to the security of ML systems. It explored

the insights of several disciplines in addition to ML, most specifically software engineering (with an emphasis on quality), security, and algorithmic game theory. It also promoted a discourse between academia and industry in a quest for well-founded practical solutions. The workshop proved to be a lively meeting of researchers from academe and industry, with presentations and discussions that were fertile and inspiring, laying the foundation for new research in the future on adversarial, reliable, and secure machine learning.

The workshop was organized by Eitan Farchi (IBM Research), Ilan Shimshoni (Haifa University), Onn Shehory (Bar Ilan University), and Anna Zamansky (Haifa University), all of whom participated in the preparation of this report. The workshop papers were published in the AAAI digital library.

Health Intelligence

Population health intelligence includes a set of activities to extract, capture, and analyze multidimensional socioeconomic, behavioral, environmental, and health data to support decision-making to improve the health of different populations. Advances in artificial intelligence tools and techniques and in internet technologies are dramatically changing the ways that scientists collect data, the ways that people interact with each other, and the ways that people interact with their environment. The internet is also increasingly used to collect, analyze, and monitor health-related reports and activities and to facilitate health-promotion programs and preventive interventions. In addition, to tackle and overcome several issues in personalized healthcare, information technology will need to evolve to improve communication, collaboration, and teamwork between patients, their families, healthcare communities, and care teams involving practitioners from different fields and specialties.

This workshop follows the success of previous health-related AAAI workshops, including those focused on personalized and population healthcare, and the first joint workshop on health intelligence at AAAI-17. This year's two-day workshop brought together a wide range of participants (roughly 60 registrants) from the multidisciplinary field of medical and health informatics. Participants were interested in the theory and practice of computational models of web-based public health intelligence, as well as personalized healthcare delivery. The full and short papers presented at the workshop covered a broad range of disciplines within AI, including knowledge representation, machine learning, natural language processing, pattern recognition, digital imaging, and online social media analytics. From an application perspective, presentations addressed topics in epidemiology, environmental and public

health informatics, disease surveillance and diagnosis, medication dosing, health behavior monitoring, and human-computer interaction.

The workshop also included three invited talks. Noémie Elhadad (Columbia University) gave a presentation on learning from patients' experiences. Nitesh Chawla (University of Notre Dame) described the value of data as a unifying component in empowering patients and personalizing their care. Amit Sheth (Wright State University) presented his work that uses AI techniques on semantically integrated multimodal data for patient-empowered health management strategies.

Martin Michalowski and Arash Shaban-Nejad served as cochairs of this workshop and coauthored this report. The workshop papers were published in the AAAI digital library.

Knowledge Extraction from Games

The Knowledge Extraction from Games workshop focused on mechanically extracting knowledge from games — including, but not limited to, game rules, character graphics, or audio environments, high-level goals or heuristic strategies, transferrable skills, aesthetic standards and conventions, or abstracted models of games. Games enjoyed by human players have been an area of interest for AI from the days of fraudulent chess automata to today's superhuman play of Go and Atari games. But games are more than just planning problems: while deep Q-learning and other efforts yield successful (albeit somewhat opaque) policies for playing specific games, we might want to ask different questions of a game system besides "How does one win?"

Games provide useful structuring information for many reasoning tasks and are therefore an ideal environment for this work. For example, games where nonplayer characters (or environment design) offer hints to solve problems might be useful stepping stones toward contextual query answering; it is not enough to find the right solution, but to identify the relationship between the textual or visual hints and the correct embodied actions. Games often share genre conventions and other similarities, or continually force a player to learn new skills or to exercise their existing competencies in novel contexts. For this reason, it seems especially interesting to explore transfer learning and analogical reasoning within and between games.

Our workshop brought together practitioners from these communities and others whose goals overlap, but whose approaches are developed in parallel — planning, general (video) game playing, knowledge representation and reasoning, knowledge extraction, computer-aided design, and others. These practitioners presented eight papers ranging

from novel cognitive architectures and transfer algorithms to new applications of image understanding and high-dimensional embeddings, from new schemes for computational creativity and knowledge transfer to inductive programming. We also hosted three invited talks from game designers (Emily Short, Raph Koster, and Ben Samuel) discussing the ways in which design knowledge is represented and reasoned about, what sorts of strategic knowledge a game might possibly contain, and how game state can effectively be communicated to players.

The key strength of this workshop was, as hoped, the integration of multiple communities of AI and automated reasoning researchers and game designers. The questions posed after the talks were stimulating, and several new collaborations were started up in the following weeks. We look forward to hosting the workshop again and to seeing the new syntheses that emerge in the next round of submissions.

The cochairs of the workshop were Joseph Osborn (University of California, Santa Cruz), Matthew Guzdial (Georgia Institute of Technology), and Adam Summerville (University of California, Santa Cruz). Joseph Osborn prepared this report. The workshop papers were published in the AAAI digital library.

Multidisciplinary Workshop on Advances in Preference Handling

Preferences are an important aspect of decision-making, and they also play an increasingly important role in computer systems. As preferences are fundamental for the analysis of human choice behavior, they are becoming of increasing importance for computational fields such as artificial intelligence, databases, and human-computer interaction, as well as for their respective applications. The Advances in Preference Handling workshop, now in its 11th iteration, provides a forum for presenting advances in preference handling and for exchanging the experiences of researchers facing similar questions, but coming from different fields.

The attendance for the day was around 25 people, and it was a lively and active group. The day was structured into four sessions: preferences in matching, markets, and games; applications; voting and preferences; and finally preference models, reasoning, and learning. There was also a panel associated with the applications session, which consisted of a one-hour discussion about the future of preference research and applications. The members of the panel were John Bresina (NASA Ames Research Center), Judy Goldsmith (University of Kentucky), Jorg Rothe (Heinrich-Heine University Dusseldorf), and Xudong Liu (University of North Florida).

The first and third sessions of the day each included five presentations on modeling and working with preferences in a variety of game-theoretic domains, including matching, markets, voting, and hedonic games. The general purpose of modeling preferences in these domains is to attempt to analytically gain insight into how agents will behave if we assume they will act in self-interested ways. The particular domains of interest — matching, markets, voting, and games — represent three fundamental domains where preferences play a key role. In each of these domains, algorithms are attempting to allocate scarce resources and at the same time be responsive to the wants and desires of the agents. Consequently, the study of preferences in these fields has assumed a large role, as the form and function of these preferences determine the efficiency and complexity of the resulting algorithms.

The late morning included an applications talk from several members of the NASA Ames research center team who are using preferences to help with satellite scheduling. This talk dovetailed nicely with the panel discussion that focused on the future of preferences and how to grow the community. Much of the discussion during the panel included ideas for new models of preferences, as well as ideas about how to better reach out to the preference learning and machine learning communities.

The last session of the day was focused on preferences models and formalisms, as well as reasoning and learning with these formalisms. These presentations all focused on classic models, including CP-nets and lexicographic preferences. Each of these presentations attempted to extend the preference modeler's toolkit with additional resources such as reasoning algorithms and distance metrics. Future iterations of the workshop will include more of this type of important and fundamental research.

Markus Endres (University of Augsburg), Nicholas Mattei (IBM Research), Andreas Pfandler (TU Wien and University of Siegen), and K. Brent Venable (Tulane University and IHMC) served as the cochairs of the workshop. Nicholas Mattei and K. Brent Venable prepared this report. Presentations were not published in the AAI digital library.

Plan, Activity, and Intent Recognition

Plan recognition, activity recognition, and intent recognition all involve making inferences about other actors from observations of their behavior, that is, their interaction with the environment and with each other. The observed actors may be software agents, robots, or humans. This synergistic area of research becomes more and more prevalent as the AI community tackles the concept of explainability, and tries to model the human-in-the-loop. This

workshop takes on this challenge by combining and unifying techniques from user modeling, machine vision, intelligent user interfaces, human or computer interaction, autonomous and multiagent systems, natural language understanding, and machine learning.

This year's workshop was centered around application domains. It included a general call regarding comparison of the various representations common in the literature and applications. We intend to respond to this call next year by providing a tutorial and other engaging activities to the community.

During this workshop, we had the pleasure to host three great speakers. Maria Gini (University of Minnesota) explored open challenges and opportunities in voice-activated intelligent personal assistants. Shirin Sohrabi (IBM Research) gave an overview of the work at IBM Research in applying plan recognition as a planning technique in several applications. Sohrabi's talk focused on the IBM Scenario Planning Advisor (SPA) tool, which is a decision-support system that utilizes plan recognition as a planning technique to assist financial organizations in identifying and managing emerging risks. Philip Cohen (Voicebox Technologies) argued that current approaches toward collaborative dialogues are too simplistic and will not extend to realistic dialogues. He instead advocated building a (joint) belief-desire-intention architecture that attempts to recognize the user's plans, and determines obstacles to their success.

The workshop was organized by Reuth Mirsky (Ben Gurion University) Sarah Keren (Technion University), and Christopher Geib (Sift LLC), who also coauthored this report. The workshop papers were published in the AAI digital library.

Planning and Inference

Planning and inference are core tasks in AI, and the connections between them have been long recognized. However, much of the work in these subareas is disjoint. The last decade has seen many exciting developments with explicit constructions and reductions between planning and inference that aim for efficient algorithms for large-scale problems and applications. The work in this area is distributed across many conferences, subcommunities, and subtopics, and includes variant flavors so that it is hard for an outsider to see a coherent picture. The goal of this workshop was to bring together researchers from these areas and to facilitate the synergy and exchange of ideas. Three topical foci were on equivalence and reductions between the problems, scalability and approximate solutions, and applications.

The workshop program included keynote talks, invited presentations, and contributed papers, and brought together diverse contributions. While the connections between the fields are obvious at a high

level, there are many ways to relate them and their algorithms at the technical level, and this was evident in the workshop.

Several papers focused on the tight links between planning and marginal MAP queries in graphical models. Several effective algorithmic approaches were introduced, including search-based inference algorithms to capture planning and influence diagrams, dual formulations of variational inference that solve influence diagrams, symbolic planning-graph variants of belief propagation, and solutions based on reductions to stochastic SAT problems. Several of these approaches used similar reductions to show how planning algorithms can solve challenging inference problems, illustrating that advances can flow in both directions.

Several talks highlighted difficulties and opportunities that arise in multiagent systems and risk-sensitive systems. Algorithmic solutions include nonconvex optimization of the inference problem, policy gradients (which is known to be closely related to the EM solution of the corresponding inference problem), and inference for beliefs over the belief of other agents.

A talk on planning in robotics showed that inference-based planning can be useful at the task level but highlighted challenges when tackling physical manipulation problems that require insight from geometry.

Two talks highlighted probabilistic programming as a framework that encompasses both planning and inference, and makes for a useful modeling and algorithmic paradigm. While bringing a distinct angle to the overall problem, at the technical level solutions use inference using variational bounds and belief modeling reminiscent of work mentioned above.

A lively discussion concluded the workshop, where participants identified challenges for future work, the potential benefits of fragmentation (for example, across discrete and continuous optimization), but on the other hand the benefits of unification since similar techniques are useful across several areas and in both fields.

Roni Khardon, Akshat Kumar, and Alex Ihler served as cochairs of the workshop. Roni Khardon prepared this report. The workshop papers were published in the AAAI digital library.

Reasoning and Learning for Human-Machine Dialogues

Natural conversation has been a key subarea of AI for decades. Its most recent form, chatbots, which can engage people in natural conversation and are easy to build, have lately been newsworthy. There is a mad rush by companies to release chatbots that can demonstrate their AI capabilities and help them gain market valuation. However, beyond basic demonstration, there is little experience in designing and

using chatbots for real-world applications that require decision-making under constraints (for example, sequential decision-making). Furthermore, there is growing interest in and a need for innovation in human-technology interaction, not least of all in the context of companion technology and social robots. Here, the aim is to implement technical systems that smartly adapt their functionality to the individual needs and requirements of their human users, even to the extent of being able to solve problems in close cooperation with those users. To this end, these AI systems need to be able to enter into dialogue and convincingly explain their suggestions and decision-making behavior.

Although statistical and machine learning methods for understanding language and detecting entities are well established, the wider problem of dialogue management remains unaddressed, as mainstream tools tend to support only rudimentary rule-based processing. The role that methods of reasoning such as constraints satisfaction, planning, and scheduling play in conversation is crucial, as is learning to work with these methods to build an end-to-end conversational system that evolves over time. On the practical side, conversational systems need to be designed to work with people such that they can explain their reasoning, persuade humans with respect to making choices, and stand up to ethical demands in real-world settings.

To address these issues, the workshop brought together over 100 AI researchers from around the world to discuss research topics centered around human-machine dialogue. The program included four invited talks, seven reviewed full paper presentations, four lightning talks accompanied by posters, and a topical panel discussion.

The day started with an invited talk by Paul Crook of Facebook. Crook argued that since there is no single type of conversation, a single approach to handle all classes of dialogue does not look possible. He explained the four approaches for building dialogue management systems — finite-state based, frame based, inference based, and response-generation based — and delved into their relative trade-offs. Erik Mueller (Capital One) then spoke about Eno, a chatbot deployed at a major bank. In the third invited talk, Matt Davis (IBM Research) explained an agent framework called CHIA, used to develop dialogue systems by building blocks of services, actions, skills, and agents. The framework has been used to quickly develop chatbots for a number of applications that are scalable and extensible, with new capabilities added over time. The final invited talk was by Kristina Jokinen (Helsinki University and NAIST), who explored interactions with social robots, as contrasted with virtual agents, detailing the issues and challenges involved in combining knowledge and dialogue capabilities for digital companions. These social robots must also take the physical state of the

humans they are planning and interacting with into account as they engage with them.

The program had the authors of peer-reviewed papers discussing ideas ranging from using neural networks, knowledge graphs, and discourse trees to generate conversation, to their applications in scheduling meetings, working within the financial industry, and learning when to speak. The discussion was wide ranging and generated many questions. The day ended with an engaging panel, moderated by Kartik Talamadupula (IBM Research) and including all the invited speakers, on the challenges inherent in wide-scale adoption of chatbots to help those in need (children, the elderly, or the physically challenged). Talamadupula noted that chatbots are equally relevant for and important to those segments of society traditionally not targeted. Further, building default behavioral assumptions into a chatbot may create issues or missed opportunities, such as a preference in some populations for short conversation or using informal language. Panelists debated the nature and category of the disadvantaged and the pros and cons of having a single agent for all or personalized agents for each, along with associated costs.

Biplav Srivastava, Susanne Biundo, Ullas Nambiar, and Imed Zitouni served as cochairs of the workshop. This report was prepared by Biplav Srivastava. The workshop papers were published in the AAAI digital library.

Statistical Modeling of Natural Software Corpora

The Statistical Modeling of Natural Software Corpora workshop was cochaired by Prem Devanbu (University of California, Davis) and William Cohen (Carnegie-Mellon University). The organizers did not provide a summary of the workshop for *AI Magazine*. The workshop papers were included in the AAAI digital library.

Bruno Bouchard is an associate professor at the Université du Québec à Chicoutimi.

Kevin Bouchard is an assistant professor at the Université du Québec à Chicoutimi.

Noam Brown is a doctoral student in computer science at Carnegie Mellon University.

Niyati Chhaya is a computer scientist in the Big Data Experience Lab at Adobe Research, Bangalore.

Eitan Farchi is a researcher at IBM Research, Haifa Research Lab.

Sebastien Gaboury is an associate professor at the Université du Québec à Chicoutimi.

Christopher Geib is a principal researcher at Smart Information Flow Technologies.

Amelie Gyraud is a postdoc researcher at Kno.e.sis, Wright State University.

Kokil Jaidka is a postdoctoral research fellow at the University of Pennsylvania.

Sarah Keren is a postdoctoral researcher at Harvard University.

Roni Khardon is a professor in computer science at Tufts University.

Parisa Kordjamshidi is an assistant professor at Tulane University.

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

Nicholas Mattei is a research staff member at IBM Research, TJ Watson Research Center.

Martin Michalowski is an assistant professor at the University of Minnesota School of Nursing.

Reuth Mirsky is a PhD candidate in the ISE Department at Ben Gurion University.

Joseph Osborn is an assistant professor in the Computer Science Department at Pomona College.

Cem Sahin is a technical staff member at MIT Lincoln Laboratory.

Onn Shehory is a professor at Bar Ilan University.

Arash Shaban-Nejad is an assistant professor at the University of Tennessee Health Science Center.

Amit Sheth is a computer scientist at Wright State University and executive director of Kno.e.sis.

Ilan Shimshoni is a professor at the University of Haifa.

Howie Shrobe is a research professor at the Massachusetts Institute of Technology.

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

Atanu R. Sinha is a principal scientist at Adobe Research, Bangalore.

Biplav Srivastava is a research staff member in the AI Department at IBM Research, Yorktown Heights.

William Streilein is a group leader at MIT Lincoln Laboratory.

Georgios Theodorou is a senior research scientist at Adobe Research, San Jose.

K. Brent Venable is an associate professor at Tulane University and IHMC.

Neal Wagner is a technical staff member at MIT Lincoln Laboratory.

Anna Zamansky is a senior lecturer in the Information Systems Department at the University of Haifa.