The AAAI-13 Conference Workshops

Vikas Agrawal, Christopher Archibald, Mehul Bhatt, Hung Bui, Diane J. Cook, Juan Cortés, Christopher Geib, Vibhav Gogate, Hans W. Guesgen, Dietmar Jannach, Michael Johanson, Kristian Kersting, George Konidaris, Lars Kotthoff, Martin Michalowski, Sriraam Natarajan, Barry O'Sullivan, Marc Pickett, Vedran Podobnik, David Poole, Lokendra Shastri, Amarda Shehu, Gita Sukthankar

■ *The AAAI-13 Workshop Program, a part of the 27th AAAI Conference on Artificial Intelligence, was held Sunday and Monday, July 14–15, 2013, at the Hyatt Regency Bellevue Hotel in Bellevue, Washington, USA. The program included 12 workshops covering a wide range of topics in artificial intelligence, including Activity Context-Aware System Architectures (WS-13-05); Artificial Intelligence and Robotics Methods in Computational Biology (WS-13-06); Combining Constraint Solving with Mining and Learning (WS-13-07); Computer Poker and Imperfect Information (WS-13-08); Expanding the Boundaries of Health Informatics Using Artificial Intelligence (WS-13-09); Intelligent Robotic Systems (WS-13-10); Intelligent Techniques for Web Personalization and Recommendation (WS-13-11); Learning Rich Representations from Low-Level Sensors (WS-13-12); Plan, Activity, and Intent Recognition (WS-13-13); Space, Time, and Ambient Intelligence (WS-13-14); Trading Agent Design and Analysis (WS-13-15); and Statistical Relational Artificial Intelligence (WS-13-16).*

Activity Context-Aware System Architectures

The third workshop on activity context-aware system architectures sought to explore architectures for intelligent context-aware systems delivering complex functionality, with direct access to information, simplifying business processes and activities while providing domain-specific and task-specific depth in interactive banking, insurance, wealth management, finance, clinical, legal, telecom customer service, operations, supply chain, connected living room, and personal assistants. Such systems understand not just words, but intentions and the context of the interaction. Such architectures are expected to dramatically improve the quality of proactive decision support provided by virtual agents by enabling them to seek explanations, make predictions, generate and test hypotheses, and perform what-if analyses. The systems can provide extreme personalization ($N = 1$) by inferring user intent, making relevant suggestions, maintaining context, carrying out cost-benefit analysis from multiple perspectives, finding similar cases from organizational and personal episodic memory before such cases are searched for, finding relevant documents and answers, and finding issues resolved by experts in similar situations. The architectures are expected to enable meaningfully relating, finding, and connecting people and information sources through discovery of causal, temporal, and spatial relations. This workshop sought to bring together researchers from the AI and human-computer interaction (HCI) communities to address key research challenges needed to create activity context-aware digital workspaces in the near future.

Key highlights of the workshops included an introduction and review by Pankaj Mehra (Fusion-io) and strong motivation for context-aware systems, followed by a keynote talk by

  ISSN 0738-4602

Benjamin Grosof (Coherent Knowledge Systems) on representing activity context through semantic rule methods, deriving from experience in the Halo and Digital Aristotle work. The afternoon session saw a keynote talk by Paul Lukowicz (DFKI) on architectures needed for large-scale collaborative social sensing. A special mention must be made of Alex Memory's (Science Applications International Corporation) architecture for context-aware insider threat detection in security applications. Genoveva G. Heredero (Infosys) demonstrated multiple applications of a detailed architecture for intelligent speech-enabled context-aware systems. The discussions in this workshop have helped us to work with fellow researchers to define activity context-aware system architectures along the following themes, which will continue to be further explored in future workshops at AAAI:

Activity Modeling, Representation, Recognition, Detection, Acquisition, Simulation and Prediction: Which (low-level) human activities can be reliably learned and detected? How indicative are those for human tasks and intent? Which granularities of activities could be chosen for creating an extensible hierarchy of human activity? What representation frameworks are most suitable for modeling activities and context switching and for enabling uniform context recall universally (across devices, platforms, and technologies)? Do we need sublanguages for user-device-specific activity and context dialogue? What types of, and how much, context information can be captured and incorporated into activity models? How can we do so effectively and efficiently?

Semantic Activity Reasoning: How to model and represent activities, objects, resources, actions, and their semantics in their context during task performance? How do we design activity/context models to enable the searching of repositories of previous activities that have behaviorally and semantically similar components to current activity requirements? What is the role of semantic memory and episodic memory? What are the techniques needed to create, manage, and properly retrieve episodic memory and to derive generalizations from episodic memory to create semantic memory, using a combination of semantic memory and episodic memory to guide users? What is the role of activity context working memory and its relationship to persistent episodic and semantic memories?

Fast Scalable Hybrid Any-Time Reasoning Systems for Context-Aware Assistance that combine numerical (and subsymbolic) and knowledge-driven (symbolic) approaches for reasoning, together with abductive reasoning, to create meaningful real-time guidance engines.

Context Information Exchange and Integration: How can we integrate and exploit the growing amount of information available from devices, services, the environment, and the various sources of general background knowledge, in order to support activity context-recognition tasks? What common ontologies or data vocabularies will be useful? What communication techniques and formalisms will be most effective in specific domains? How can the externalized cognitive state transfer be properly affected? What are the relevant use-case scenarios and collaboration environments? What are suitable software architectures, user interfaces, developer tools, and benchmarking tools for activity-based computing? What kind of text, context, and behavioral analytics are needed?

Pankaj Mehra, Lokendra Shastri, and Vikas Agrawal organized this workshop. This summary was written by Vikas Agrawal and Lokendra Shastri. The papers of the workshop were published as AAAI Press Technical Report WS-13-05.

Artificial Intelligence and Robotics Methods in Computational Biology

In the last two decades, many computer scientists in artificial intelligence and robotics have made significant contributions to modeling biological systems. Indeed, the fields of computational structural biology are now highly populated by researchers with diverse backgrounds in search, planning, learning, evolutionary computation, constraint programming, machine learning, data mining, and others, and great progress is being made on methods to solve problems related to structure prediction, motion simulation, and design of biological macromolecules (proteins and nucleic acids).

The workshop brought together a diverse group of researchers from a variety of subfields within AI and robotics, such as robot motion planning, evolutionary computation, machine learning, computational geometry, and kinematics. Three keynote speakers summarized seminal progress in biomolecular modeling. A talk given by Mona Singh (Princeton University, Department of Computer Science) highlighted research on understanding and modeling molecular interactions with a variety of techniques from machine learning and computational geometry. A talk by Pierre Baldi (University of California Irvine, School of Information and Computer Sciences) outlined the historical developments and the general principles behind deep architectures for learning, and showcased powerful applications for prediction of biomolecular structure. A talk by Yang Zhang (University of Michigan, Department of Computational Medicine and Bioinformatics) focused on prediction of tertiary structure in proteins, highlighting state of the art research on threading-based, template-free, and hybrid approaches and important lessons and open questions from successes and failures in CASP.

Workshop presentations addressed five major themes. Papers on the computation of molecular motions highlighted robotics-inspired algorithms for modeling folding paths in proteins and detailed mapping of structural transitions in peptides (by Juan Cortés, Lydia Tapia, and Nancy Amato). A second group of papers highlighted recent developments on evolutionary search algorithms and the particular promise of multiobjective optimization for modeling equilibrium structures of loops and polypeptide chains in proteins (by Amarda Shehu and Yaohang Li). Another theme was assisted protein structure prediction with prior information on contacts (by Jianlin Cheng). Advanced mathematical concepts were introduced to enhance protein structure determination from data obtained by X-ray crystallography by exploiting

symmetry (by Gregory Chirikjian). A global view of the relationship between sequence, structure, and function in the protein universe was also provided (by Rachael Kolodny).

Amarda Shehu, Juan Cortés, and Jianlin Cheng cochaired this workshop. This summary was written by Amarda Shehu and Juan Cortés. The papers of the workshop were published as AAAI Press Technical Report WS-13-06.

Combining Constraint Solving with Mining and Learning

The field of constraint solving has traditionally evolved quite independently from the fields of machine learning and data mining. In recent years, interest has grown around the opportunities at the interfaces between these fields, and the potential advantages of their integration. Integration can work in two ways; on the one hand, various types of constraint solvers can be included in machine-learning and data-mining algorithms, for example, to provide a uniform and effective way to characterize the desired solutions. On the other hand, machine learning can help address challenges in constraint satisfaction and programming, both at the level of search, by improving search or integrating intelligent metaheuristics, as well as at the level of modeling, for example, by learning constraints or interactively supporting a decision maker.

While promising initial results have been achieved at these interfaces, many opportunities remain unexplored and further research is needed in order to establish a systematic approach to this integration. The main purpose of this workshop was to provide an open environment where researchers in machine learning, data mining, and constraint solving could exchange ideas and discuss promising approaches, crucial issues, open problems, and interesting formal- izations of new tasks.

The workshop program comprised two invited talks and six peer-reviewed research papers. Both invited talks presented interesting perspectives on the use of constraint reasoning to support machine-learning tasks. The first invited talk was given by Dan Roth (University of Illinois at Urbana-Champaign, USA) on amortized integer linear programming inference. The second talk was given by Ian Davidson (University of California, Davis, USA) on incorporating constraint satisfaction into hierarchy building, graph segmenting, and clustering. Three of the peer-reviewed papers addressed topics related to the use of machine learning to assist constraint solving and applications, specifically, to improve search (Loth et al.), predict algorithm runtime (Hutter et al.), and facilitate interesting complex transportation problems (Aleksandrov et al.). The remaining three papers considered the use of constraints in various machine-learning and data-mining contexts, specifically, using constraints to assist the learning of Bayesian learning (Yao et al.), feature construction (Costa), and using constraints for pattern mining (Guns et al.). Therefore, the program was quite evenly split between papers from either a constraints or a machine-learning background.

The mix of attendees from the key scientific topic areas of the workshop was a significant feature of the workshop and was commented on throughout the event. There were many shared interests among the participants, and it was felt that this event, the second in the series, provided a basis for an interesting and productive annual meeting for people with complementary interests and skills in these important research areas. A number of collaborations have been initiated as a result of the workshop.

The workshop cochairs are grateful to all the contributors to the workshop for helping make the event such a success, and to the program committee members who reviewed every submission and selected the papers that were presented. We are also grateful to the *Artificial Intelligence* journal for financial support.

Tias Guns (KU Leuven, Belgium), Lars Kotthoff (University College Cork, Ireland), Barry O'Sullivan (University College Cork, Ireland), and Andrea Passerini (University of Trento, Italy) cochaired this workshop. This report was written by Lars Kotthoff and Barry O'Sullivan. The papers of the workshop were published as AAAI Press Technical Report WS-13-07.

Computer Poker and Imperfect Information

Poker is the canonical game with imperfect information and stochastic events. Agents playing the game must attempt to maximize their winnings against an initially unknown adversary in a large and complex environment, while trying to discover what hidden information the opponent knows and acting in order to hide their own private information. The task of creating agents that can match or surpass professional players in these human-scale domains is an exciting and active area of research.

Since 2006, the Annual Computer Poker Competition (ACPC) has used the game of poker as a common test bed for researchers investigating the application of artificial intelligence to stochastic imperfect information domains. Significant progress has been made over the last eight years of the competition: computers have surpassed human players at two-player limit Texas Hold'em (the smallest game played for high stakes by humans) and are now approaching optimal play, and computer agents are continuing to improve in the more complex no-limit and three-player variants. The Computer Poker and Imperfect Information Workshop brings together scientists and competitors from the ACPC with researchers who study the broad class of imperfect information domains.

A wide range of techniques from artificial intelligence and machine learning are used in poker and other imperfect information domains, and the workshop provides a venue for researchers to present their research as oral and poster presentations. A continuing theme, this year and historically, is the progress made in the field of computational game theory, where competitors attempt to approximate optimal Nash equilibrium strategies that are guaranteed to do no worse than tie in two-player zero-sum games. The ACPC has resulted in the development of two algorithms that are more efficient in both time and memory than previous approaches for solving large imperfect information games. Presentations in this area focused on state-space approximation techniques, effi-

cient implementations of the state-of-the-art counterfactual regret minimization algorithm, and techniques for solving the game as a series of subproblems.

Another major theme this year was the use of online and offline techniques for modeling an agent's adversaries and varying the agent's actions to better respond to them. This directly addresses the goal of the game of poker, which is to maximize the expected winnings against an initially unknown group of adversaries, but historically has been a more challenging problem than precomputing a robust nonadaptive agent. There was active debate on defining the objective of the task, a presentation of a new online adaptation technique used in the ACPC, and applications of machine-learning techniques to building offline models of agent behavior.

The results of the 2013 Annual Computer Poker Competition were also announced at the workshop and can be found at www.computerpokercompetition.org. The ACPC is run over the two months leading up to the workshop with the goal of producing statistically significant results between every combination of competitors in every event.

Finally, the workshop was used to discuss the past and future of the Annual Computer Poker Workshop. This began with an eight-year retrospective presentation by Michael Johanson, who discussed the research progress that the community has made over time. This set the stage for a discussion of the future directions of the ACPC and the workshop, including proposed rule changes, the introduction of new events, and challenges and opportunities facing the research area as a whole.

Christopher Archibald and Michael Johanson cochaired this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-13-08.

Expanding the Boundaries of Health Informatics Using Artificial Intelligence

The rapid development and expansion of health informatics results in a constantly growing volume of available data with diversified structure and format that is coupled with an increased reliance on this data by health practitioners and patients. The main boundaries to fully leveraging the data are its sophisticated management, analysis, and use of results. AI techniques are very well suited to expand these boundaries, thus facilitating advances in the health informatics area that may have a profound effect on patient outcomes. As such, the Health Informatics Using Artificial Intelligence workshop is a forum to bring together health informatics researchers and AI researchers with the goal of improving patient outcomes through innovation.

This year's workshop received a large number of submissions that were divided into two major tracks (application and methodology related). In addition to these tracks, one keynote and two invited speakers provided crucial insights into and directions for health informatics research. The keynote presentation, Artificial Intelligence in Medicine: It's Back to the Future given by Mark Musen (Stanford University), captured the theme of the workshop. In his talk, Musen described the current opportunities for innovation in health informatics thanks to advances in technology and changes in how medicine is delivered. He specifically advocated opportunities for resuscitating AI in health informatics and the need for venues such as HIAI to learn from past lessons and to exchange ideas to foster future discoveries.

The application track began with an invited talk by Jay M. Tenenbaum (Cancer Commons) discussing opportunities to apply human and machine intelligence to organize and refine the knowledge about treating cancer. The papers in this track talked about the use of ontologies to define, organize, integrate, and process information in patient health records and population health records. Additional insights into the use of semantic networks to describe and execute computational experiments helped round out the track.

The theory track began with an invited talk by Barry O'Sullivan (University College Cork) on the challenges and opportunities for applying constraint programming to gaining health knowledge. Due to the diversity of presented topics, this track was split into two sections. The first one grouped topics related to classification and prediction problems, while the second dealt with other complex problems. In the first section, papers focused on the construction of classification and prediction models using machine-learning techniques on different types of clinical data (images, typical alphanumeric information, and textual reports). The second section encapsulated work on the analysis of other types of clinical decision problems with a focus on identifying interactions and addressing them between concurrently applied CPGs, the strategies to preemptive assignment of caregivers to patients during mass casualty incidents, and the analysis of patient reviews of doctors.

To further foster interaction and help guide future research direction, the invited speakers were joined by Wojtek Michalowski (University of Ottawa) to carry on a panel discussion. The panelists, with participation from the audience, discussed a number of topics including physicians' trust in AI's ability to help them make decisions, AI researchers' responsibility in physician decisions aided by AI-based support systems, and the need to train physicians in understanding complex decision models. There was general agreement that physicians should have some understanding of the technology helping them make decisions. Further, for the proposed AI solutions to be accepted in practice we as researchers need to provide solutions to existing needs even if these problems are low level and simple to implement.

Martin Michalowski, Wojtek Michalowski, Dympna O'Sullivan, and Szymon Wilk cochaired this workshop. This report was written by Martin Michalowski. The papers of the workshop were published as AAAI Press Technical Report WS-13-09.

Intelligent Robotic Systems

Robotics has shown dramatic progress in recent years, driven in part by the development of standardized hardware and open source software frameworks that dramatically lower the effort required to obtain a working robot system. These advances represent an ex-

citing opportunity for AI researchers, because they have brought robot technology to the point where there is a pressing need for integrating AI techniques into robot systems. The task of designing complete, intelligent robotic systems presents us with the opportunity to develop fully fledged agents that interact with the real world, and the challenge of coping with the complexity and uncertainty that such interaction entails, leading to an immensely rich source of research directions that have the potential for significant real-world impact. Robotics also offers us an opportunity to integrate the disparate subfields of AI, and drive progress in each subfield through the grand challenge of intelligent robotics.

This workshop built on the success of two prior events, Designing Intelligent Robots: Reintegrating AI, at the AAAI Spring Symposia at Stanford, and the AI Meets Robotics meetings in Örebro and Lyon. The workshop drew participants from all over the world, including five invited speakers and 20 contributed papers, and closed with a lively discussion on future directions.

Byron Boots, Nick Hawes, Todd Hester, George Konidaris, Tekin Merili, Lorenzo Riano, Benjamin Rosman, and Peter Stone cochaired this workshop. This report was written by George Konidaris. The papers of the workshop were published as AAAI Press Technical Report WS-13-10.

Intelligent Techniques for Web Personalization and Recommendation

The role of personalization of web-based systems and the automatic recommendation of potential items of interest to users is probably higher today than ever before: search engines personalize more and more their results for a particular user, social media sites automatically highlight or filter activity feeds and suggest friends to connect to, and companies even announce million dollar prizes for improvements in the area of e-commerce personalization. Web personalization in general aims to tailor the web experience to a particular user or set of users. The goals of personalization can be comparably simple, for example, making the presentation more pleasing, but can also be complex when the aim, for example, is to anticipate the needs of a user and provide information in a customized form. Recommender systems represent one special and prominent class of personalized web applications, which focus on the user-dependent filtering and selection of relevant information and aim to support online users in the decision-making and buying process. The recent developments in the area of recommender systems — in particular in the context of the social web — generate new demands, in particular with respect to interactivity, adaptivity, and user-preference elicitation. These challenges, however, are also in the focus of general web personalization research.

The workshop therefore aimed to bring together practitioners and researchers from the partially overlapping fields of web personalization and recommender systems to discuss current and emerging topics in their respective fields and to foster an exchange of ideas and experiences.

The workshop was opened with an invited talk by Guy Shani from the Information Systems Engineering department at the Ben Gurion University, Israel. In his talk, Shani reported on recent results on user interaction aspects in recommender systems applications and in particular on the role of relevance displays and how they affect the user acceptance of such systems.

The technical papers presented in different themed sessions were selected in a peer-review process by an international program committee and covered a variety of topics related to web personalization and recommendation. The topics of the papers ranged from the automated generation of music playlists and revenue-maximizing movie recommendations, over approaches for interactive search and configuration to social and contextual recommendation techniques.

The discussions after the technical paper presentations were centered on questions related to the use of contextual information in the recommendation process as well as on problems of designing and conducting user studies for evaluation purposes.

Dietmar Jannach, Sarabjot Singh Anand, and Bamshad Mobasher cochaired this workshop. This report was written by Dietmar Jannach. The papers of the workshop were published as AAAI Press Technical Report WS-13-11.

Learning Rich Representations from Low-Level Sensors

A human-level artificially intelligent agent must be able to represent and reason about the world, at some level, in terms of high-level concepts such as entities and relations. The problem of acquiring these rich high-level representations has long been an obstacle for achieving human-level AI. A popular approach to this problem is to handcraft these high-level representations, but this has had limited success. An alternate approach is for rich representations to be learned autonomously from low-level sensor data. Potentially, the latter approach may yield more robust representations and should require less reliance on human knowledge engineering.

The goal of this workshop was to provide an informal forum for discussing sensor-learning approaches to the problem of how a machine may acquire a broad range of knowledge about the world. Approximately 25 participants attended this workshop to discuss the 11 accepted papers and 4 invited talks representing different approaches to bridging the gap between low-level sensors and rich high-level representations. Papers included approaches from areas such as developmental psychology, neuroscience, deep learning, robotics, machine vision, and reinforcement learning.

Although the approaches differed, some recurring themes emerged from the discussion. One theme was the importance of learning hierarchical representations, whether these are feature hierarchies in connectionist networks or hierarchies of macro-actions in reinforcement learning. In general, hierarchies were learned in a bottom-up manner, with lower-level features making it possible for systems to learn higher-level representations. Another theme was the importance of generality and modality independence of learning algorithms. If a large amount

of knowledge is to be learned, then learning algorithms that are tailored to a specific modality (such as vision) won't be as flexible as more general-purpose learning mechanisms.

The workshop also raised several questions. One open question is whether distinct algorithms are required for learning different levels of representation, or whether it is possible for a single algorithm to learn all these levels. Jeff Hawkins described a system, modeled on the human cortex, that was essentially the same general-purpose learning and inference algorithm repeated many times. Alternatively, Benjamin Kuipers and Joseph Modayil both described systems that spanned much of the bridge from raw sensors to high-level representations, but that used different algorithms for different stages. Kuipers outlined a "40-year-vision" on how the bridge could be completely spanned.

Another question is how a system can learn relations from sensor data that isn't explicitly relational. For example, how can the relation "above" be learned given only the raw pixels of many still images? This is related to the question of how causality might be learned from sensor data. Amy Fire demonstrated a system that was able to learn perceptual causality from videos. Both George Konidaris and Jonathan Mugan made headway into the question of how symbols might be learned from sensor data. Konidaris presented a system that creates symbols in a continuous reinforcement learning domain, while Mugan described a system that learns qualitative representations from a continuous domain in a top-down manner.

Finally, both Yoshua Bengio and Juergen Schmidhuber presented impressive recent advances from the deep learning community. For example, Schmidhuber presented a system that learns to drive a simulated car given video data. An open question is how deep learning can be extended to learn relational representations, such as those used in natural language.

At the end of the workshop, participants engaged in a general discussion. Although no definite conclusions were reached, the importance of learning representations for achieving human-level AI was stressed, as well as the need for researchers from disparate areas to compare notes on this common obstacle to these areas.

Marc Pickett, Benjamin Kuipers, Yann LeCun, and Clayton Morrison cochaired this workshop. The report was written by Marc Pickett. The papers of the workshop were published as AAAI Press Technical Report WS-13-12.

Plan, Activity, and Intent Recognition

Plan, activity, and intent recognition (PAIR) all involve making inferences about other actors from observations of their behavior, for example, their interaction with the environment or with each other. Techniques for plan, activity, and intent recognition play a crucial role in a wide variety of applications including personal assistant technology, security systems, gaming and simulation, human-robot interaction, and automated dialogue systems. For this reason, the problem has continued to receive attention from researchers in a number of areas: human-computer interaction, autonomous and multiagent systems, machine vision, and natural language understanding.

An important topic concerning the PAIR community is how to make efficient and accurate inference about an agent's hidden plans and intents from the potentially noisy observation of the agent's external behavior. Following the tradition of PAIR, many technical presentations at the workshop continue to push the envelope in addressing this key issue. Approaches presented include constructing expressive probabilistic-logical representations of an agent's behavior (Song, Kautz, and Luo), dealing with agents acting in exploratory manner (Uzan, Dekel, and Gal), strategies for parallelization (Geib and Swetenham), incremental abduction (Meadows, Langley, and Emery). Also presented was empirical analysis on how well different techniques work in a number of different domains, including activity recognition benchmarks (Ross and Kalleher; Nazerfard and Cook) and mobile text prediction (Freedman et al.). Taking a more general perspective in the context of multiagent interaction, a cluster of presentations at the workshop this year illustrated the interesting synergy between models for planning (taking action) and plan recognition (interpretation) (Khan, Arif, and Boloni; Smith and Lieberman; Vered and Kaminka).

This year, we were delighted to have Prashant Doshi (University of Georgia) and Jerry Hobbs (USC ISI) as our guest speakers; their invited talks were one of the highlights of the workshop. The new poster session provided additional opportunities for workshop participants to have more in-depth technical discussions at the side of the posters. The workshop concluded with a traditional group dinner, which many of the participants attended.

Hung Bui, Gita Sukthankar, Christopher Geib, and David V. Pynadath cochaired this workshop. The report was written by Hung Bui, Gita Sukthankar, and Christopher Geib. The papers of the workshop were published as AAAI Press Technical Report WS-13-13.

Space, Time, and Ambient Intelligence

The Space, Time, and Ambient Intelligence workshops focus on basic research questions concerned with the computational modeling of commonsense situational awareness for assistive technologies within the purview of ambient intelligence and smart environments. Space, time, and ambient intelligence research topics address systems concerned with observing, interpreting, and interacting in an environment populated by humans and artifacts; the emphasis is on formal methods for representing and reasoning about spatiotemporal-, event-, and action-driven phenomena that occur in the environment or domain of interest.

This workshop focused on the topic of spatiotemporal aspects of human activity interpretation, welcoming research concerned with the interpretation of people interactions, real-time commonsense situational awareness involving aspects such as scene perception and understanding, perceptual data analytics, and prediction and explanation-driven high-level control of autonomous systems. In addition to six paper presentations, the workshop featured two invited keynote speakers: An-

thony Cohn (University of Leeds, UK) and Henry Kautz (University of Rochester, USA).

The technical program of the workshop also included prototypical demonstrations and initiatives on benchmarking and promotion of open-access algorithms and systems from the viewpoint of cognitive vision, interaction, and control. The two discussion sessions at the workshop also focused on benchmarking and tool development with an emphasis on commonsense abstraction and reasoning (about human activities) encompassing aspects such as common sense, space, and change. Some questions that were discussed included what is the status quo in open-source development and benchmarking in high-level cognitive interpretation? What kind of interfacing (for example, logic programming interface, APIs) and syntactic sugar would be needed to use commonsense methods for modeling, reasoning, and learning about space, actions, events, change, and interaction? How could reasoning plug in to large-scale or hybrid activity recognition / interpretation projects? What kinds of general developmental, debugging, visualisation capabilities would be needed to support a broad range of cognitive vision projects?

A general consensus was that research and development in the field needs to work on open availability and accessibility of data (for example, RGB and depth profile), and tool sets for spatiotemporal abstraction, reasoning, learning, and data visualisation.

Mehul Bhatt, Hans W. Guesgen, and Diane J. Cook cochaired this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-13-14.

Trading Agent Design and Analysis

Trading agents are a prominent area of research in artificial intelligence and multiagent systems. The design and analysis of trading agents poses significant challenges in decision making, and many different artificial intelligence techniques have been combined in the study of trading agents, including planning, decision theory, game theory, machine learning, optimization, and others. In addition to research interest, trading agents have potential benefits in electronic commerce, supply-chain management, and other business interests. Since 2003 Trading Agent Design and Analysis (TADA) workshops have been focused on all aspects of the design and evaluation of trading agents, including agent architectures, decision-making algorithms, theoretic analysis of agents or market games, empirical studies of agent performance, agent negotiation strategies, game-theoretic studies, market architectures, and other related topics.

Additionally, a special characteristic of Trading Agent Design and Analysis workshops is a set of papers related to Trading Agent Competition scenarios, covering analyses of strategies used in previous Trading Agent Competitions, discussions about the effectiveness of different approaches, as well as thoughts on applying the lessons learned through the Trading Agent Competition to other domains. Trading Agent Design and Analysis has a long and successful history being colocated at different artificial intelligence conferences, including two previous workshops hosted by AAAI in 2007 and 2008.

The international program committee comprising 26 researchers from academe and industry selected seven papers for presentation at the workshop. Papers presented at the workshop spanned from empirical analyses of real-world Internet-based trading agent systems to initial explorations of trading agent solutions in emerging application areas such as smart grid environments. The workshop started with two papers dealing with auctions: auctioneer profitability in QuiBids penny auctions and price dependencies in simultaneous sealed-bid auctions were analyzed. The remainder of the presentations were related to the Trading Agent Competition: there was a paper on detection of opportunistic bids in the TAC Supply Chain Management scenario as well as a set of papers from a domain of the new Trading Agent Competition scenario focusing on energy markets. Power Trading Agent Competition papers covered challenges such as electricity demand forecasting using Gaussian processes, smart charging for electric vehicles using reinforcement learning, prediction of customer demand in energy market simulations using machine learning, as well as analysis of Power Trading Agent Competition trading agent key performance indicators. The workshop concluded with a demonstration of the AdX game, a new version of a Trading Agent Competition game in the domain of Internet ad auctions and a discussion on emerging application areas of trading agents.

The Trading Agent Design and Analysis workshop traditionally coincides with the final rounds of the Trading Agent Competition and many of the contestants participate in the workshop venue. The Trading Agent Competition finals are also used as a way to promote research on trading agents to a broader audience and to generate interest in new application areas for agent technologies. For example, the Power Trading Agent Competition, a new Trading Agent Competition scenario focusing on smart grid markets, generated significant interest within the community and has attracted new workshop participants with interests in energy applications. The inaugural edition of Power Trading Agent Competition finals as well as the Trading Agent Competition Ad Auctions finals were collocated with the TADA-13 workshop. The TacTex team from the University of Texas at Austin (USA) led by Daniel Urieli won the first Power Trading Agent Competition, while the TAU agent from Tel Aviv University (Israel), led by Mariano Schain, performed best in the Trading Agent Competition Ad Auctions.

The Trading Agent Competition board members provided guidance on organization matters, and Ioannis Vetsikas (NCSR Demokritos, Greece) was involved in organizational matters related to the Trading Agent Competition.

Vedran Podobnik chaired this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-13-15.

Statistical Relational Artificial Intelligence

Much has been achieved in the field of AI, "the science and engineering of making intelligent machines" as John

McCarthy defined it, yet much remains to be done if we are to reach the goals we all imagine. One of the key challenges with moving ahead is closing the gap between logical and statistical AI. Logical AI has mainly focused on complex representations, and statistical AI on uncertainty. Clearly, however, intelligent machines must be able to handle the complexity and uncertainty of the real world.

Recent years have seen an explosion of successes in combining probability and (subsets of) first-order logic, programming languages, and databases in several subfields of AI such as reasoning, learning, knowledge representation, planning, databases, natural language processing, robotics, vision, and others.

Nowadays, we can learn probabilistic relational models automatically from millions of interrelated objects. We can generate optimal plans and learn to act optimally in uncertain environments involving millions of objects and relations among them. Exploiting shared factors can speed up message-passing algorithms for relational inference but also for classical propositional inference such as solving SAT problems. We can even perform lifted probabilistic inference, avoiding explicit state enumeration by manipulating first-order state representations directly. So far, however, the researchers combining logic and probability in each of these subfields have been working mostly independently. This workshop was designed for attempts at synthesis, forming a common core of problems and ideas, cross-pollinating across subareas, and ultimately starting to explore what might be called statistical relational AI: the study and design of intelligent agents that act in noisy worlds composed of objects and relations among the objects.

The 21 papers and posters at the workshop covered a wide range of statistical relational AI topics such as lifted inference, natural language processing, online rule learning, event recognition, tractable relational representations, relational Markov decision processes (MDPs), and learning statistical relational learning (SRL) models, thus clearly showing the promise of statistical relational (SR) techniques for AI.

In addition, the workshop had three invited speakers, Dan Suciu (University of Washington), Prasad Tadepalli (Oregon State University), and Toby Walsh (NICTA). The topics of these presentations explored the connections between statistical relational AI community and (1) probabilistic databases, (2) constraint satisfication, and (3) relational MDPs. Another interesting part of the workshop was a lively poster session that encouraged the participants to discuss the commonalities and need for differences among the various AI tasks that can be addressed by SR techniques. The group reached a general consensus that statistical relational AI is an exciting emerging area requiring more investigation. The topic of efficient and lifted inference found particular interest.

Kristian Kersting, Vibhav Gogate, Sriraam Natarajan, and David Poole cochaired this workshop and wrote this report. The papers of the workshop were published as AAAI Press Technical Report WS-13-16.

Vikas Agrawal is a research staff member at IBM Research-India.

Christopher Archibald is an assistant professor in the Department of Computer Science and Engineering at Mississippi State University.

Mehul Bhatt is a researcher in the Cognitive Systems, and Spatial Cognition Research Center at the University of Bremen, Germany.

Hung Bui is a principal research scientist at the Laboratory for Natural Language Understanding, Nuance, Sunnyvale, CA.

Diane J. Cook is a chaired professor in the School of Electrical Engineering and Computer Science at Washington State University.

Juan Cortés is a researcher with the LAAS-CRNS Laboratory at the University of Toulouse.

Christopher Geib is an associate professor in the College of Information Science and Technology, Drexel University.

Vibhav Gogate is an assistant professor at the Department of Computer Science, University of Texas at Dallas, USA.

Hans W. Guesgen is a professor of computer science in the School of Engineering and Advanced Technology at Massey University, New Zealand.

Dietmar Jannach is a chaired professor of computer science at TU Dortmund, Germany.

Michael Johanson is a Ph.D. student at the University of Alberta in Edmonton, Canada.

Kristian Kersting heads the statistical relational mining group at Fraunhofer IAIS and the University of Bonn, Germany.

George Konidaris is a postdoctoral associate at MIT.

Lars Kotthoff is a postdoctoral research scientist at the INSIGHT Centre for Data Analytics, University College Cork.

Martin Michalowski is a principal research scientist at Adventium Labs in Minneapolis, MN.

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

Barry O'Sullivan is a full professor in the Department of Computer Science at University College Cork (Ireland) and director of the INSIGHT Centre for Data Analytics, University College Cork.

Marc Pickett is an NRC postdoctoral fellow at the Naval Research Laboratory in Washington, DC.

Vedran Podobnik is an assistant professor at the Telecommunication Department of the Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia.

David Poole is a professor at the Department of Computer Science, University of British Columbia, Vancouver, Canada.

Lokendra Shastri is the associate vice president, GM Research, and head of the Center for Knowledge Driven Intelligent Systems at Infosys, India.

Amarda Shehu is an assistant professor in the Department of Computer Science at George Mason University.

Gita Sukthankar is a Charles N. McMillan faculty fellow and an associate professor in the Department of Electrical Engineering and Computer Science at the University of Central Florida.