

Reports of the AAAI 2010 Conference Workshops

*David W. Aha, Mark Boddy, Vadim Bulitko,
Artur S. d'Avila Garcez, Prashant Doshi,
Stefan Edelkamp, Christopher Geib,
Piotr Gmytrasiewicz, Robert P. Goldman,
Alon Halevy, Pascal Hitzler, Charles Isbell,
Darsana Josyula, Leslie Pack Kaelbling,
Kristian Kersting, Maithilee Kunda, Luis C. Lamb,
Bhaskara Marthi, Keith McGregor, Lilyana Mihalkova,
Vivi Nastase, Sriraam Natarajan, Gregory Provan,
Anita Raja, Ashwin Ram, Mark Riedl, Stuart Russell,
Ashish Sabharwal, Jan-Georg Smaus, Gita Sukthankar,
Karl Tuyls, and Ron van der Meyden*

■ The AAAI-10 workshop program was held Sunday and Monday, July 11–12, 2010, at the Westin Peachtree Plaza in Atlanta, Georgia. The AAAI-10 workshop program included 13 workshops covering a wide range of topics in artificial intelligence. The titles of the workshops were *AI and Fun*; *Bridging the Gap Between Task and Motion Planning*; *Collaboratively Built Knowledge Sources and Artificial Intelligence*; *Goal-Directed Autonomy*; *Intelligent Security*; *Interactive Decision Theory and Game Theory*; *Metacognition for Robust Social Systems*; *Model Checking and Artificial Intelligence*; *Neural-Symbolic Learning and Reasoning*; *Plan, Activity, and Intent Recognition*; *Statistical Relational AI*; *Visual Representations and Reasoning*; and *Abstraction, Reformulation, and Approximation*. This article presents short summaries of those events.

AI and Fun

Interactive entertainment has become a dominant force in the entertainment sector of the global economy. In 2000, John Laird and Michael van Lent justified interactive entertainment as a domain of study in AI when they posited that computer games could act as test beds for achieving human-level intelligence in computers, leveraging the fidelity of their simulations of real-world dynamics. There is an additional perspective on AI for games: increasing the engagement and enjoyment of the player. This perspective is consistent with the perspective of computer game developers. For them, AI is a tool in the arsenal of the game to be used in lieu of real people when no one is available for a given role. Examples of such roles are opponents, companions, and nonplay characters (NPCs) in roles that are not fun to play such as shopkeepers, farmers, and victims; cinematographer; dungeon master; plot writer; or game designer.

As we move down this list, the computational system is charged with progressively more responsibility for providing a user with a fun experience.

But what is fun? We seem to know it when we see it, but fun is also highly subjective. Can we computationally model fun? Can intelligent systems learn and utilize models of fun, player preferences, storytelling, and so on, to affect human experiences? If so, what would this enable with respect to increased engagement, enjoyment, or new forms of computer-mediated interactive experiences? What are the potential ways forward? To begin to address these questions, we invited 11 research groups to present on their work and to help shed light on the questions from a number of perspectives. We explicitly did not ask for papers, but instead asked that each group present on their perspective on AI approaches to fun, to address the questions above, and to speculate on the future. The presentations roughly clustered into four themes: (1) Player modeling and learning from humans, featuring presentations on how to learn models of players, customize game play experience, and make inferences about what humans like to do. (2) Virtual and real humans, featuring presentations on virtual humans in mixed-reality game environments, modeling human improvisational actors for the purposes of building better interactive experiences, and agents that express curiosity. (3) Storytelling and discourse, featuring presentations on the question of whether we can achieve the dream of the Holodeck, how interactive storytelling might reinforce cognitive perception of engagement, and how computational systems can mediate virtual experiences through automated cinematography. (4) Making learning AI fun, featuring presentations on the question of how to engage learners in AI courses with virtual worlds and robotics.

In addition to the four themes, we invited three experts from a disparate set of industries to talk about challenges and opportunities for AI and fun: Miguel Encarnacao, director of emerging technology innovation at Humana, Inc.; Joe Marks, vice president of Disney Research; and Bob Sottolare, chief technology officer of the U.S. Army Simulation and Training Technology Center. While it would seem that there is little in common with training warfighters, making people healthy, and entertaining people in theme parks, a consensus emerged that fun is important for motivating people to learn, proactively maintain their health, and create brand loyalty. There was also a call for more research into natural forms of human-computer communication. Most notably, all three domains cited scalability—more people, more personalization, longer experience durations, longitudinal interactions—as a primary bottleneck that required creative automated intelligence.

While we are not ready to formally define the term *fun*, we can—and should—use it as a call to arms for an investigation into core research on intelligent systems that reason about and manage the quality of human experiences both in a variety of domains—including many beyond games, such as education, training, and health—and in a variety of computer-mediated experiences such as storytelling, interactive drama and theater, serious games, mixed-reality environments and virtual cinematography. While fun can be subjective, progress can be made through study of related, objectively measurable phenomena: engagement, enjoyment, immersion, flow, replayability, motivation, and others yet to be identified. Finally, we note the strong potential for societal impact through the domains that can be affected as well as insight into the basic questions of what drives us and how we can computationally model and automatically reason about it.

Mark Riedl, Charles Isbell, Ashwin Ram, and Vadim Bulitko served as cochairs of this workshop. No technical report was issued.

Bridging the Gap Between Task and Motion Planning

Task-level planning of the sort typically studied by the AI and planning communities has historically been quite separate from motion planning as studied by roboticists. There is an increasing belief that it is both necessary and useful for a tighter coupling between these levels. From the perspective of task planning, motion planning can be viewed as providing geometric constraints and heuristics, and therefore information about the feasibility and costs of higher-level actions. From a motion planning perspective, task planners allow taking advantage of the rich combinatorial structure that exists in the configuration spaces that arise when, for example, manipulating several objects.

The workshop featured talks from diverse areas. There were several presentations on particular formalisms for combining the levels of planning from the perspectives of classical planning, motion planning, POMDPs, search-based planning, and temporal logic-based constraints. There were also presentations on robotic applications and open source software, as well as fast GPU implementations of motion planners.

Several common themes emerged from the presentations and discussion. There was a general consensus that pure low-level motion planning in configuration space was fairly well solved for problems of moderate dimensionality, and that a principled approach to coupling with higher-level constraints and goals was the next step. There were a variety of opinions, however, on the form that this coupling should take, ranging from hierarchical

approaches that tackle the different levels in concert, to approaches based on finding a good discretization or summary of lower-level state, but then planning independently. Finally, many participants stressed the need for common benchmarks and representations, so that the different approaches could be compared.

The workshop was organized by Maxim Likhachev, Bhaskara Marthi, Conor McGann, and David E. Smith. The papers of the workshop were published as AAAI Technical Report WS-10-01.

Collaboratively Built Knowledge Sources and Artificial Intelligence

Until recently, the AI, and in particular the natural language processing (NLP), communities have relied on resources built manually by experts in specific areas (such as linguists, philosophers, and cognitive linguists). User-contributed knowledge has opened up a new perspective, in that it captures the kind of knowledge and organization that arises naturally out of the consensus of the masses, and as such represents better our collective knowledge. The outcome is a multifaceted and extremely rich source of information, revealed through embedded annotations and structural information.

The first such collaboratively developed repository of information to be extensively used in AI and natural language processing was Wikipedia. Its usefulness was demonstrated through its contributions to a wide range of tasks: text categorization, clustering, word-sense disambiguation, information retrieval, information extraction, and question answering.

In recent years, more and more resources and collaborative endeavours have started to be incorporated and exploited as knowledge repositories for various tasks. Tags associated with images in Flickr, question-answer collections in Yahoo! Answers are a few examples of such information sources. Amazon's Mechanical Turk gives researchers access to "human computation" power, and is being used more and more as a solution to the difficult problems of large-scale evaluations and data annotation, both crucial for the continuous development of the AI and NLP fields.

The Collaboratively Built Knowledge Sources and Artificial Intelligence workshop took place on July 11, 2010, in Atlanta, Georgia, immediately preceding the 24th AAAI Conference on Artificial Intelligence. This workshop is a successor to the Wikipedia and Artificial Intelligence: An Evolving Synergy workshop organized at AAAI 2008 and the User Contributed Knowledge and Artificial Intelligence: An Evolving Synergy workshop organized as part of IJCAI 2009.

Consistent with the original aim of the work-

shop, the presented papers address a diverse set of problems and resources. Some papers explored the process and result of combining different resources, extracting and formalizing knowledge from structured, semistructured or unstructured sources, harnessing people power and learning from them how to automatically generate and enhance knowledge repositories. The two invited talks explored in more detail two of the workshop themes. Henry Lieberman (Massachusetts Institute of Technology Media Laboratory) presented work developing a platform for knowledge acquisition and usage from expert users, built upon experience with eliciting information from nonexperts in a collaborative manner, and Lenhart Schubert (University of Rochester) explored the usage of knowledge extracted from general text for formal inference, and the impact of ambiguity and specific constructs in language on the knowledge extracted and the inferencing methods.

Vivi Nastase, Roberto Navigli, and Fei Wu served as cochairs of this event. Supporting agencies included AAAI, and HITS, gGmbH, Heidelberg, Germany. The papers of the workshop were published as AAAI Technical Report WS-10-02.

Goal-Directed Autonomy

The objective of the AAAI Workshop on Goal-Directed Autonomy (GDA) was to encourage discussion and novel contributions on intelligent agents that can self-select their goals. How should an autonomous agent behave competently when interacting in a complex environment (for example, partially observable, multiagent, with large decision spaces, dynamic updates, stochastic outcomes, and continuous effects)? The option of complete a priori domain engineering is not appealing due to its high cost; it would require planning for all possible contingencies due to opportunities or plan execution failures. Alternatively, the agent could be given the ability to decide what goals it should pursue at any point in time, which would increase its level of autonomy by relaxing the assumption that its assigned goal is the only one that it should pursue throughout its lifetime. This capability of goal reasoning could dramatically affect the types of tasks that these agents can perform.

As demonstrated by the workshop's attendees, goal reasoning is of interest to researchers studying cognitive architectures, game AI, multiagent systems, planning, robotics, and other topics in which competent agent behavior is desirable in complex environments. To our knowledge, this was the first workshop that focused on goal reasoning. We named the workshop GDA because it is the name of a recent conceptual model for goal reasoning whose components address problems

relevant to other such models. These include (1) detecting situations that may trigger goal reasoning, (2) explaining why a detected situation demands attention, (3) deciding how to respond to such situations (for example, through goal(s) formulation), and (4) managing the current set of pending goals, which may involve tasks such as goal interruption, transformation, resumption, and/or deletion.

The workshop began with a survey on goal reasoning, given by Matthew Klenk (NRL). Felipe Meneguzzi (Carnegie Mellon University) discussed work on motivations, and Matt Molineaux (Knexus Research) discussed progress on a recently evaluated system that builds on this foundation for goal formulation. In the context of cognitive architectures and meta reasoning, Dongkyu Choi (Stanford University) discussed how goal reasoning is embedded in the ICARUS, Randy Jones (SoarTech) described the use of appraisal theory for implementing GDA in Soar, and Ashok Goel (Georgia Institute of Technology) described its relation to metareasoning.

In a robotics context, Jeremy Baxter's (QinetiQ) presentation focused on user interaction, while Nick Hawes (University of Birmingham) described a goal-directed reasoning framework for controlling a mobile robot. Héctor Muñoz-Avila (Lehigh University) and Ben Weber (University of California, Santa Cruz) described different roles of case-based reasoning for GDA, and Russell Knight surveyed applications of CASPER, a real-time embedded planner scheduler. Finally, Mike Cox (DARPA) monitored a panel on the relation of goal reasoning to plan adaptation (Muñoz-Avila), replanning (Ugur Kuter, University of Maryland), and planning and uncertainty (Daniel Bryce, Utah State University).

There were many interruptions and lively discussion on topics such as the approaches for designing the components of a GDA model (for example, when, how, and what new goals should be formulated?), the relative benefits of alternative models, how to manage concurrent goals, the relation of goal reasoning to automated subgoaling and opportunistic planning, and the ability of current methods to scale.

David Aha, Matthew Klenk, Héctor Muñoz-Avila, Ashwin Ram (Georgia Institute of Technology), and Daniel Shapiro (ISLE) served as organizers of this workshop. The papers presented were not published.

Intelligent Security

The purpose of the Intelligent Security workshop series was to bring together researchers with an interest in both security and AI. The goals were to tease out common themes and differences, identify common problems and their solutions, share

experiences with the applicability of techniques from one field to problems from the other, and to identify the key issues to be addressed in increasing the convergence between security and AI.

This is a fertile area for research, and has been attracting an increasing amount of interest in both communities. Prior to this workshop there was a 2009 ICAPS workshop on the topic, as well as two workshops held in conjunction with the ACM Conference on Computer and Communications Security (CCS), and so organized primarily from the computer security community.

AI and security is a large and growing area, both for research and for applications. Our increasingly networked world continues to provide new opportunities for security breaches that have severe consequences at the personal level (identity theft, and resulting financial losses), for businesses (theft of intellectual property, or business plans, or costly responses to the theft of customer data), and for governments. Computing and the Internet have become crucial parts of the infrastructure of almost every significant commercial or governmental enterprise. Turning off the computers or disconnecting from the network has become tantamount to turning off the power.

The use of techniques drawn from AI is increasingly relevant as the scale of the problem increases, in terms of the size and complexity of the networks being protected, in terms of the variety of applications and services provided using that infrastructure, and with the sophistication of the attacks being made. Filtering the faint signals of intrusion from a flood of data related to normal operations can be viewed as data mining. Learning methods can be applied to generate classifiers for this process, or to detect the presence of new means of attack. AI planning methods can be used to generate compact representations of possible attacks, which can then be used to deploy countermeasures. Plan and intent recognition are important areas of research as well and are the focus of a growing number of researchers. The detection of anomalous operations or network traffic can be viewed as a component of many security functions, including both intrusion detection and plan recognition. Another recent topic is improving anomaly detection using the ubiquitous and increasingly powerful graphics processors in our computers. Because of the distributed nature of computer networks, they are susceptible to attack that comes from multiple directions, which can be mounted by an individual in a single location. Thus, the issue of information fusion (combining indications drawn from separate data streams) is an important tool, as well.

Mark Boddy, Stefan Edelkamp, and Robert P. Godman served as cochairs of this AAAI workshop. The papers presented were not published.

Interactive Decision Theory and Game Theory

The Interactive Decision Theory and Game Theory workshop is a continuation of a series of workshops on decision and game theories held over previous years. These topics remain active research areas since game and decision theories proved to be powerful tools with which to design autonomous agents and to understand interactions in systems composed of many such agents. Decision theory provides a general paradigm for designing agents that can operate in complex uncertain environments and can act rationally to maximize their preferences. Decision-theoretic models use precise mathematical formalism to define the properties of the agent's environment, the agent's sensory capabilities, the ways in which the agent's actions change the state of the environment, and the agent's goals and preferences. Agent's rationality is defined as behavior that maximizes the expectation of the degree to which the preferences are achieved over time, and the planning problem is identified as a search for the rational, or optimal, plan.

Game theory adds to the decision-theoretic framework the idea of multiple agents interacting within a common environment. It provides ways to specify how agents, separately or jointly, can change the environment and how the resulting changes affect their individual preferences. Building on the assumption that agents are rational and self-interested, game theory uses the notion of Nash equilibrium to design mechanisms and protocols for various forms of interaction and communication that result in the overall system behaving in a stable, efficient, and fair manner.

Recent research has sought to merge advances in decision and game theories to build agents that may operate in complex uncertain environments shared with other agents. This research has investigated the problems of Nash equilibrium as a solution concept, focused on epistemological advances in game theory and expressive ways to model agents. Alternative solution concepts have been investigated with the aim of designing autonomous agents that robustly interact with other, highly sophisticated agents in both cooperative and noncooperative settings.

Papers presented at the workshop spanned the spectrum of theoretical issues as well as emerging application areas. There were papers on learning to cooperate, computation of steady states in two-player extensive games, improved fast computation of Nash equilibria, and the maximum entropy approach to computing correlated equilibria. Papers that included applications were ones on cognitive hierarchies applied to the lemonade game, updating higher order beliefs during bar-

gaining, opponent modeling and Monte-Carlo search in poker, trust models applied in supply chain management, and a computational decision-theoretic approach for interactive assistants. Still other papers were devoted to teamwork and coordination under model uncertainty and to signaling games as models for conversational grounding. Overall, we were impressed with the technical maturity and high level of formalism presented in the papers, and the impressive selection of applications the researchers are looking at.

Piotr Gmytrasiewicz, Prashant Doshi, and Karl Tuyls served as cochairs of this event. The papers of the workshop were published as AAAI Technical Report WS-10-03.

Metacognition for Robust Social Systems

The one-day Workshop on Metacognition for Robust Social Systems was a sequel to a series of successful workshops on the topic of metareasoning beginning in 2007. The focus of this workshop was on design considerations, issues, and challenges in using metacognition to improve the robustness of social systems that include purely artificial entities or both humans and software agents. The workshop had both full paper and four-page position papers, some of which built on results from the previous workshops—an encouraging sign of the formation of a research community in this area. The papers were categorized under three themes: (1) metacognition in human-machine social systems, which included discussions on modeling human behavior, reproducing humanlike interactions, understanding and dealing with conflicts in human input, and the role of modeling self and emotions in human-machine social systems; (2) metacognition in multiagent systems, including domains suitable for decentralized metacognition and a common platform for evaluating multiagent cognition to the implementation of decentralized metacognition in multiagent systems; and finally (3) metacognitive architectures, a theme that included classification of metacognition under higher-level, theory-based metacognition and lower-level, experience-based metacognition, metarepresentational theories of metacognition applied for theory-based metacognition, control theories of metacognition applied for experience-based metacognition and frameworks for implementing scalable metacognitive architectures.

In addition to the presentations, the sessions were followed by an interactive panel where each author presented views and comments regarding topics put forth by a selected moderator. These sessions also gave the audience an opportunity to ask questions or make comments on issues that

spanned individual paper presentations. Some of the issues discussed were the complexities and issues involved in modeling agents situated in a social system that includes humans (versus all machines); the role of emotions in social systems; the feasibility of having models of self and models of others; criteria for evaluating metacognitive systems; deadline with overhead costs associated with metareasoning; and finally modeling emotions to potentially help with metacognitive processing.

A special highlight of the workshop was the invited talk. Ashok Goel of the Georgia Institute of Technology spoke of model-based metareasoning for self-adaptation. He first presented the idea of proactive, goal-directed reconfiguration of reasoning processes. He used an assembly task as a motivating example to discuss how reasoning mechanisms of one task can be analogically transferred to other tasks. Professor Goel then discussed localized reinforcement learning in the context of free-civ, a turn-based strategy game and showed how model-based metareasoning would provide reward signals for different parts of the search space. He concluded the talk with his thoughts on retrospective, failure-driven repair of domain knowledge. This was followed by a lively question and answer session on the role of models in metareasoning.

The cochairs of this workshop were Anita Raja (University of North Carolina at Charlotte) and Darsana Josyula (Bowie State University). The papers were published as AAI Technical Report WS-10-04.

Model Checking and Artificial Intelligence

Model checking is an approach to verification based on representing some system as a model, a semantic structure that supports the truth or falsity of a formula of a logic. Typically, the model describes all the possible behaviors of a system over time, and the logic is a modal logic that allows one to specify some desired or undesired behaviors in a concise way. A model checker is a software system that takes as inputs representations of the system and its specification in modal logic, and computes whether the specification holds in the system.

The interactions between model checking and artificial intelligence are rich and diverse. Model checking originated in the 1980s as an approach to the verification of concurrent hardware processes and computer network communications protocols. These days, model checking is applied by researchers working on computer software such as hardware device drivers and cryptographic protocols. AI applications include planning, stochastic process models, autonomous robots, and other forms of multiagent systems. This broadening of application area has also led to a broadening of the

range of modal logics, and there are now model checkers whose specification language is able to express modalities such as knowledge, belief, and so on, as well as time. Such modalities are of particular interest when dealing with autonomous and multiagent systems.

In principle, a model checker conducts an exhaustive examination of the state space. Underpinning the success of the area is a range of sophisticated optimization techniques and heuristic algorithms that enable this computation to be performed efficiently. In this regard, model checking has benefited from a range of ideas from artificial intelligence, including search heuristics such as A^* , and planning approaches to counterexample construction.

Several themes were touched by the papers presented at the workshop. Stefan Leue (University of Konstanz) presented an algorithm for finding the k shortest paths in graph, a problem that is relevant, among others, for stochastic model checking. Stefan Edelkamp (University of Bremen) spoke on the use of the graphical processing unit (GPU) for external memory breadth-first search. Hector Geffner (University of Barcelona) gave an invited talk on planning with incomplete information, stressing that while logic and theory are needed in planning, the bottom line is heavily empirical. Two further presentations were also devoted to planning. Siddharth Srivastava (University of Massachusetts) spoke on computing applicability conditions for plans with loops, with various results concerning termination and other behaviours of transition systems applying not just to a particular planning formalism, and hence of interest to the model checking community. Stefan Edelkamp spoke on action planning for automated program verification, using approaches from planning for verification of C code.

Multiagent systems and epistemic logic was another workshop theme, building a strong bridge between AI concepts and model checking. Abdur Rakib (University of Nottingham) spoke on bounded model checking for temporal epistemic logic. Francesco Russo (Imperial College, London) spoke on automatic data abstraction in model checking multiagent systems. Ron van der Meyden presented an extension of CTL with epistemic operators. Finally, Kaile Su (University of Beijing and University of Brisbane) gave a presentation on the Herbi-vore protocol, which involves knowledge and anonymity.

During the workshop discussion period, Stefan Edelkamp aimed at bridging the gap between planning and model checking and raised the question of whether the knowledge aspect is really essential for the applications just mentioned.

Not surprisingly, the workshop was geared more toward AI than model checking, and there was

some support of the idea that to reach into the model-checking community, the next workshops should move away from AI conferences for a change.

Ron van der Meyden and Jan-Georg Smaus served as coauthors of this workshop. The papers of this workshop were not published.

Neural-Symbolic Learning and Reasoning

AI faces huge challenges in its quest to develop truly intelligent systems. These systems are required to learn and adapt to changes in the environment they operate, and to reason about commonsense knowledge in ways that can control the accumulation of errors. Current developments in the area of neural-symbolic computation bring an opportunity to integrate well-founded symbolic AI reasoning and inference systems with robust neural computing machinery and learning to help tackle some of these challenges.

Neural-symbolic systems combine the statistical nature of learning and the logical nature of reasoning. Over the years, researchers have built sound neural-symbolic models that are able to learn several forms of reasoning, including temporal, modal, epistemic, fuzzy, intuitionistic, and relational (first-order, predicate) logics. In a nutshell, neural-symbolic computation offers a methodology for integrating reasoning and learning in intelligent systems and a rich model for cognitive computation. These features allow the integrated study of symbolic and connectionist AI. Further, neural-symbolic computation seeks to provide explanations to certain important questions in cognitive science, such as the nature of reasoning, knowledge representation, and learning, following the computational theory of mind.

The workshop contained a mix of invited talks and submitted presentations. Invited talks spanned the foundations of (logical) reasoning and neural and statistical learning to the application of neurosymbolic technology in the aerospace industry, training simulators, and vehicle control.

The workshop started with a keynote address by philosopher Paul Thagard who presented exciting new ideas on how brains make mental models. Thagard brought the role of abductive reasoning to the foreground and discussed the importance of attention and creativity to this reasoning task, with particular emphasis on how emotions can drive one's attentional focus.

Gadi Pinkas gave an invited talk on how to represent first-order logic in symmetric networks, particularly Boltzmann machines. The effective integration of first-order logic and artificial neural networks has been a challenge for decades. Pinkas revisited his proposed model and discussed recent

The Feigenbaum Prize

The first biennial AAAI Feigenbaum Prize will be awarded in 2011 at the 25th anniversary of the AAAI Conference in San Francisco, California. The AAAI Feigenbaum Prize recognizes and encourages outstanding Artificial Intelligence research advances that are made by using experimental methods of computer science. The "laboratories" for the experimental work are real-world domains, and the power of the research results are demonstrated in those domains.

The Feigenbaum Prize may be given for a sustained record of high-impact seminal contributions to experimental AI research; or it may be given to reward singular remarkable innovation and achievement in experimental AI research.

The prize is \$10,000 and is provided by the Feigenbaum Nii Foundation and administered by AAAI.

For complete details about how to submit nominations for this prize, please see www.aaai.org/Awards/feigenbaum.php

developments, implementations and challenges, including complexity and first-order learning issues.

Leo de Penning presented a paper introducing an integrated neural-symbolic cognitive agent architecture for training, assessment, and feedback in simulators. De Penning illustrated the practical use of neural-symbolic computation in a virtual instruction platform, reporting initial results on a real simulator environment. The proposed architecture combines temporal logic reasoning and learning in recurrent restricted Boltzmann machines.

Sihle Wilson presented a paper that contains the initial design of a neuro-fuzzy plug-in that would allow vehicles autonomously to retrieve a driver from a nearby location. The vehicle would use neural networks to help it avoid collisions, but also fuzzy logic to help it make suboptimal decisions in case a collision cannot be avoided.

Dragos Margineantu gave an invited talk reporting on years of experience at Boeing and DARPA on testing adaptive and learning models. Margineantu presented a robust methodology for evaluating learning systems. Particular emphasis was placed on the evaluation of high-risk low-probability events, the difficulties associated with small

or zero counts, and the challenges relating to high-cost anomaly detection.

Pedro Domingos's talk advocated the use of a foundational approach to the research in the area. Various methods used by the logical and statistical approaches to AI were reviewed, with emphasis on Markov logic networks (MLNs) as a proposed synthesis toward a unified theory of learning and cognition. MLNs integrate probabilistic inference and first-order logic in novel ways and have been shown useful in a range of applications.

Kristian Kersting's invited talk on statistical relational mining discussed the exponential growth of scientific data and exciting new directions on efficient large-scale inference and knowledge extraction.

Jim Prentzas and Ioannis Hatzilygeroudis presented a paper that tackles the problem of knowledge extraction from neural networks (and knowledge extraction from graphical models or complex networks, in general) — a crucial part of NeSy systems but also an important research area in its own right.

The discussion session was coordinated by Artur d'Ávila Garcez. D'Ávila Garcez listed a large number of challenges and areas for further research. The discussion then focused on two main issues: endowing agent systems with emotions and representations for learning, including nonclassical logic and deep networks. Other issues included handling inconsistency, model checking and the modeling of beliefs and intentions.

Artur S. d'Ávila Garcez, Pascal Hitzler, and Luis C. Lamb served as coauthors of this event. The papers of this workshop were not published.

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. This area of research combines and unifies techniques from user modeling, machine vision, intelligent user interfaces, human or computer interaction, autonomous and multiagent systems, natural language understanding, and machine learning. It plays a crucial role in a wide variety of applications including assistive technology, software assistants, gaming, computer and network security, behavior recognition, coordination in robots and software agents, e-commerce and collaborative filtering. This diversity of applications and disciplines, while producing a wealth of ideas and results, has contributed to fragmentation in the field, as researchers publish relevant results in a wide spectrum of journals and conferences. This work-

shop brought together researchers from these diverse areas to share ideas and developments.

The workshop included two invited talks. Hector Geffner spoke on plan recognition as planning. Tanzeem Choudhury spoke on activity recognition and social network identification using small mobile platforms. Both invited speakers, along with Gita Sukthankar, were panelists on a very interesting panel discussion of how to go about learning the structures that are to be recognized. As the speakers used very different models, this resulted in some very interesting and diverse ideas and lively debate. Among the major themes of the workshop were a more formal understanding of the limits of some approaches to plan recognition, learning plans for recognition by demonstration, and activity and behavior recognition for use in home environments.

Christopher Geib, Gita Sukthankar, David Pynadath, and Hung Bui served as coauthors of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-10-05.

Statistical Relational AI

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

The last decade has seen real progress toward closing the gap. 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 the AI subfields have been working mostly independently.

The Statistical Relational AI workshop convened researchers driving forward work in this area through the interplay between addressing AI tasks and using statistical relational techniques to solve them, forming a common core of problems and ideas, and ultimately starting to explore what we call "statistical relational AI": the science and engi-

neering of making intelligent machines that act in noisy worlds composed of objects and relations among the objects.

The 19 papers and posters at the workshop covered a wide range of topics including Bayesian abductive reasoning, lifted inference, lifted planning and planning by probabilistic programming, lifted SAT, relational learning, probabilistic programming for natural language processing, relational data integration, cognitive architectures, and killer applications, among others.

In the first part of the workshop, researchers presented new technical contributions. Bart Selman's and Josh Tenenbaum's invited talks provided a synthesis of probabilistic and logical inference and a statistical relational AI perspective on acquiring commonsense theories respectively. The second 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 statistical relational 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, Stuart Russell, Leslie Pack Kaelbling, Alon Halevy, Sriraam Natarajan, and Lilyana Mihalkova served as cochair of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-10-06.

Visual Representations and Reasoning

Visual representations and reasoning play an important role in human problem solving, modeling, and design. Although the ability to think like a human long has been a goal of AI, today's AI agents nonetheless are limited in their visual reasoning. Advances in this area may enable more extensive autonomous reasoning in visual domains, foster deeper computational support for and understanding of human problem solving, modeling, and design, and promote more intense use of visual representations in human-machine interaction. These technological goals raise basic theoretical issues such as the precise role of visual reasoning in intelligence and the relationship between visual reasoning and perceptual processes. This interdisciplinary workshop aimed to describe and discuss the latest scientific research that may inform and influence progress toward these goals.

The workshop brought together participants from diverse research communities such as AI, cognitive science, learning science, and design science, in addition to researchers in the fields of philosophy, public policy, health systems, and human-computer interaction. One major theme of

papers presented at the workshop was the primacy of spatial reasoning, especially in such nonobvious domains as blind-map navigation and the diagnosis of certain mental illnesses. Generally, these papers built upon the major theme of the AAAI 2009 fall symposium on multirepresentational architectures for human-level intelligence. The workshop also included an invited talk, given by Paul Rosenbloom (University of Southern California) on the development and research direction of SOAR 9, and its focus on the integration of first-class perceptual processing.

The workshop concluded with a lively discussion and debate among the participants concerning the place of visual and spatial reasoning in the pantheon of artificial intelligence research. While the utility of purely visual representations in traditional physical symbol systems was apparent to those in attendance, considerable discussion occurred regarding the realization of perceptual symbol systems.

Keith McGreggor and Maithilee Kunda served as cochair of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-10-07.

Abstraction, Reformulation, and Approximation

It has been recognized since the inception of artificial intelligence that abstractions, problem reformulations, and approximations (ARA) are central to human commonsense reasoning and problem solving and to the ability of systems to reason effectively in complex domains. ARA techniques have been used in a variety of problem-solving settings and application domains, primarily to overcome computational intractability by decreasing the combinatorial costs associated with searching large spaces. In addition, ARA techniques are also useful for knowledge acquisition and explanation generation.

The aim and scope of this one-day workshop were similar to an independent and bigger symposium series called SARA (Symposium on Abstraction, Reformulation, and Approximation). Since the early 1990s, the SARA symposia have provided an opportunity to bring together participants with diverse backgrounds, leading to a rich and lively exchange of ideas, allowing the comparison of goals, techniques, and paradigms, and helping identify important research issues and engineering hurdles. This workshop continued to do the same, filling in a gap that was felt in the recent years as SARA has lately been organized every other year. The workshop provided a somewhat more informal setting than the symposia, without formal published proceedings but with an AAAI Press

Technical Report documenting the papers presented at the workshop.

This was the first year for this workshop. With 11 papers grouped into three themes—ARA for constraint-satisfaction problems, ARA for machine learning, and ARA for reasoning in combinatorial search spaces—the workshop brought together an exciting range of topics to the table, especially those of interest to AAAI conference participants. The talks touched upon topics such as satisfiability solvers, knowledge compilation, state space abstraction for search, reinforcement learning, inductive logic programming, and answer set programming. The active presence of senior researchers in the area such as Robert Holte (University of Alberta, Canada) helped provide a broad perspective throughout the workshop, especially during the open discussion session held toward the end.

The participants emphasized strong interest in the general area of abstraction in computer science and indicated that having this workshop fill the gap between the 2009 and 2011 SARA symposia was highly useful in maintaining a feeling of continuity in this research community.

Gregory Provan and Ashish Sabharwal served as cochairs of this workshop. The papers of the workshop were published as AAAI Press Technical Report WS-10-08.

David W. Aha leads the Adaptive Systems Section at the Naval Research Laboratory in Washington, D.C.

Mark Boddy is a member of the technical staff at Adventium Labs, Minneapolis, Minnesota.

Vadim Bulitko is an associate professor in the Department of Computer Science, University of Alberta.

Artur S. d'Avila Garcez is a reader in computing at City University London, UK.

Prashant Doshi is an associate professor at the University of Georgia.

Stefan Edelkamp is a researcher and lecturer at the University of Bremen, Germany.

Christopher Geib is a research fellow in the School of Informatics at the University of Edinburgh.

Piotr Gmytrasiewicz is an associate professor at the University of Illinois at Chicago.

Robert P. Goldman is a principal scientist at Smart Information Flow Technologies LLC in Minneapolis, Minnesota.

Alon Halevy heads the Structured Data Group at Google Research.

Pascal Hitzler is an assistant professor of computer science at Wright State University, Dayton, Ohio.

Charles Isbell is an associate professor in the School of Interactive Computing, Georgia Institute of Technology.

Darsana Josyula is an assistant professor of computer science at Bowie State University.

Leslie Pack Kaelbling is a professor of computer science and engineering at the Massachusetts Institute of Technology.

Kristian Kersting heads the Statistical Relational Mining group at Fraunhofer IAIS and is affiliated with the Computer Science Department at the University of Bonn, Germany.

Maithilee Kunda is a computer-science doctoral student at Georgia Institute of Technology, Atlanta, Georgia.

Luis C. Lamb is an associate professor of computer science at the Federal University of Rio Grande do Sul, Brazil.

Bhaskara Marthi is a research scientist at Willow Garage, Inc., in California.

Keith McGregor is a computer-science doctoral student at Georgia Institute of Technology, Atlanta, Georgia.

Lilyana Mihalkova is a member of the LINQS group at the Computer Science Department, University of Maryland College Park.

Vivi Nastase is a research associate at the Heidelberg Institute for Theoretical Studies in Heidelberg, Germany.

Sriram Natarajan is a post-doctoral research associate at the Department of Computer Science in the University of Wisconsin Madison.

Gregory Provan is a professor in the Department of Computer Science at University College Cork, Ireland, and directs the Cork Complex Systems Lab.

Anita Raja is an associate professor of software and information systems at the University of North Carolina at Charlotte.

Ashwin Ram is an associate professor in the School of Interactive Computing, Georgia Institute of Technology.

Mark Riedl is an assistant professor in the School of Interactive Computing, Georgia Institute of Technology.

Stuart Russell is a professor of computer science and engineering at the University of California, Berkeley.

Ashish Sabharwal is a research associate in the Department of Computer Science at Cornell University, USA, and works at the Institute for Computational Sustainability.

Jan-Georg Smaus is a Privatdozent at the University of Freiburg, Germany.

Gita Sukthankar is an assistant professor in the School of Electrical Engineering and Computer Science at the University of Central Florida.

Karl Tuyls is an associate professor at Maastricht University, The Netherlands.

Ron van der Meyden is a professor at the University of New South Wales and NICTA, Australia.



Preliminary Call for Papers

The Fifth International AAAI Conference on Weblogs and Social Media

July 17–21, 2011 ☒ Barcelona, Catalonia, Spain

Sponsored by the Association for the Advancement of Artificial Intelligence

ICWSM 2011 will be held on July 17–21, 2011 in Barcelona, Catalonia, (Spain) and will be collocated with IJCAI 2011. The International AAAI Conference on Weblogs and Social Media (ICWSM) is a unique forum that brings together researchers from the disciplines in computer science, linguistics, communication, and the social sciences. The broad goal of ICWSM is to increase understanding of social media in all its incarnations. Submissions describing research that blends social science and technology are especially encouraged.

Although this conference is relatively new, it has become one of the premier venues for social scientists and technologists to gather and discuss cutting-edge research in social media. This is largely due to a typical acceptance rate of 20 percent for full-length research papers and support from the Association for the Advancement of Artificial Intelligence (AAAI).

For ICWSM 2011, in addition to the usual program of contributed technical talks, tutorials, and invited presentations, the main conference will include a selection of keynote talks from prominent social scientists and technologists. This year, for the first time, we will be expanding our workshop program and are planning an additional workshop day in addition to the usual tutorial and main conference. We have received a number of exciting workshop proposals already and are looking forward to receiving many more. Also, since we will be collocating with IJCAI-11 there will be opportunities to connect with people in that community both formally and informally (a joint social event is in the planning).

Disciplines

- Computational Linguistics/NLP
- Text Mining/Data Mining
- Psychology
- Sociology (including Social Network Analysis)
- Anthropology, Communications, Media Studies
- Visualization
- Political Science
- Computational Social Science
- HCI
- Economics
- Graph theory, concrete analysis and simulation of graphical models

Important Dates

Abstract Submission: January 31, 2011
Full Paper Submission: February 7, 2011
Notification of Acceptance: March 18, 2011
Camera Ready Due: April 4, 2011
Conference in Barcelona: July 17–20, 2011

Data Workshop Paper Submission: March 22, 2011
Workshop Paper Acceptance Notification: April 8, 2011

Tutorial Proposal Submission: February 18, 2011
Tutorial Acceptance: March 4, 2011

Workshop Proposal Submission: January 7, 2011
Workshop Acceptance: January 18, 2011
Workshop Paper Submission: March 22, 2011
Workshop Paper Acceptance Notification: April 8, 2011

Media

- Weblogs, including comments
- Social Networking Sites
- Microblogs
- Wikis (Wikipedia)
- Forums, usenet
- Community media sites: youtube, flickr

Topics

Topics include the following:

- Psychological, personality-based, and ethnographic studies of social media
- Analyzing the relationship between social media and mainstream media
- Qualitative and quantitative studies of social media
- Centrality/influence of social media publications and authors
- Ranking/relevance of blogs, web page ranking based on blogs



© David Pedraza Loureiro

The chimneys on top of La Pedrera

- Social network analysis, communities identification, expertise and authority discovery, collaborative filtering
- Trust, reputation, recommendation systems
- Human computer interaction, social media tools, navigation and visualization
- Subjectivity in textual data, sentiment analysis, polarity/opinion identification and extraction
- Text categorization, topic recognition, demographic/gender/age identification
- Trend identification and tracking, time series forecasting, measuring predictability of phenomena based on social media
- New social media applications, interfaces, interaction techniques

Keynotes

- Sinan Aral (Stern Business School, NYU)
- Manuel Castells (Annenberg School for Communication, USC)
Tentative

Submissions

People interested in participating should submit through the ICWSM-11 website a technical paper (up to 8 pages, not including references), poster or demo description (up to 4 pages) by the deadlines given above (Midnight PST). Papers must be formatted in AAAI two-column, camera-ready style (see the AAAI author instructions page at www.aaai.org/Publications/Author/author.php). Details for the submission procedure will appear at the conference website (icwsm.org) in December 2010.

Submissions to Other Conferences or Journals

ICWSM-11 will not accept any paper that, at the time of submission, is under review for or has already been published or accepted for publication in a journal or conference. This restriction does not apply to submissions for workshops and other venues with a limited audience.

Registration

All accepted papers and extended abstracts will be published in the conference proceedings. At least one author must register for the conference by the deadline for camera-ready copy submission. In addition, the registered author must attend the conference to present the paper in person.

Publication

All accepted papers and abstracts will be allocated eight (8) pages in the conference proceedings. Authors will be required to transfer copyright of their paper to AAAI.

Data Challenge

ICWSM-11 will once again hold a data challenge featuring a freely-available dataset and a half-day workshop at the conference. Details will be posted on the conference website.

Conference Website

www.icwsm.org/2011/index.shtml

General Information

For general information regarding ICWSM-11, please write to icwsm11@aaai.org.

Organizing Committee

General Chairs

Nicolas Nicolov (Microsoft)

James G. Shanahan (Independent Consultant)

Program Committee Area Chairs

Lada Adamic (University of Michigan)

Ricardo Baeza-Yates (Yahoo Research)

Scott Counts (Microsoft Research)

Local Chairs

Alex Jaimes (Yahoo! Labs)

Ricard Ruiz de Querol, (Barcelona Media Innovation Centre)

Sponsorship Chairs

Ron Kass (Pidgin Technologies)

Data Challenge Chair

Ian Soboroff

Demo Chair

Ido Guy (IBM Research)

Publicity Chairs

danah boyd (Microsoft Research)

Tutorial Chair

John Breslin (NUI Galway)

Webmaster

Ritesh Agrawal (AT&T Interactive)

Webdesigner

Disco Gerdes (Mutual Mobile)

Venue

Barcelona, Spain

Student Awards

We will be providing a limited number of student awards to help cover the cost of travel, subsistence, and registration to the ICWSM 2011 conference. Details will be posted on the conference website.

Call for Papers and Participation



The Twenty-Third IAAI Conference on Artificial Intelligence

San Francisco, California, USA 9–11 August 2011 Collocated with AAAI-11

Sponsored by the Association for the Advancement of Artificial Intelligence

The Twenty-Third Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-11) will focus on successful applications of AI technology. The conference will use technical papers, invited talks, and panel discussions to explore issues, methods, and lessons learned in the development and deployment of AI applications; and to promote an interchange of ideas between basic and applied AI.

IAAI-11 will consider papers in two tracks: (1) deployed application case studies and (2) emerging applications or methodologies. Submissions should clearly identify which track they are intended for, as the two tracks are judged on different criteria (see definitions and criteria, below). All submissions must be original.

Deployed Application Case Study Papers

Case-study papers must describe deployed applications with measurable benefits that include some aspect of AI technology. The application needs to have been in production use by its final end-users (not the people who created it) for sufficiently long that the experience in use can be meaningfully collected and reported. This period typically spans at least three months. The case study may evaluate either a stand-alone application or a component of a complex system. In addition to the criteria listed below for Emerging Track papers, the deployed applications will also be evaluated on the following:

Task or Problem Description: Describe the task the application performs or the problem it solves. State the objectives of the application and explain why an AI solution was important. If other solutions were tried and failed outline these solutions and the reasons for their failure.

Application Description: Describe the application, providing key technical details about design and implementation. What are the system components, what are their functions, and how do they interact? What languages and tools are used in the application? How is knowledge represented? What is the hardware and software environment in which the system

is deployed? Provide examples to illustrate how the system is used.

Uses of AI Technology: On what AI research results does the application depend? What key aspects of AI technology allowed the application to succeed? How were the techniques modified to fit the needs of the application? If applicable, describe how AI technology is integrated with other technology. If a commercial tool is used, explain the decision criteria used to select it. Describe any insights gained about the application of AI technology. What AI approaches or techniques were tried and did not work? Why not?

Application Use and Payoff: How long has this application been deployed? Explain how widely, how often, and by whom the application is being used. Also describe the application's payoff. What measurable benefits have resulted from its use? What additional benefits do you expect over time? What impacts has it had on the users' business processes?

Application Development and Deployment: Describe the development and deployment process. How long did they take? How many developers were involved? What were the costs? What were the difficulties, and how were they overcome? What are the lessons learned? What, if any, formal development methods were used?

Maintenance: Describe your experience with and plans for maintenance of the application. Who maintains the application? How often is update needed? Is domain knowledge expected to change over time? How does the design of the application facilitate update?

Original papers on the deployment issues in AI applications are welcome, even if other papers on the AI technology have been presented at or submitted to other conferences. We encourage updates on applications that have been in use for an extended period of time (i.e., multiple years). Each of the accepted deployed

application papers will receive the IAAI "Innovative Application" Award.

A copy of the review form for Deployed Application Case Study Papers will be posted on the IAAI website (www.aaai.org/Conferences/IAAI/iaai11.php) in the fall of 2010.

Emerging Application Case Study Papers

The goal of the emerging application track is to "bridge the gap" between basic AI research and deployed AI applications, by discussing efforts to apply AI tools, techniques, or methods to real world problems. Emerging applications are on aspects of AI applications that are not appropriate for deployed application case studies, or are not sufficiently deployed to be submitted as case studies. This track is distinguished from reports of scientific AI research appropriate for the AAAI-11 Conference in that the objective of the efforts reported here should be the engineering of AI applications.

Emerging application papers may include any aspects of the technology, engineering, or deployment of AI applications, including discussions of prototype applications; performance evaluation of AI applications; ongoing efforts to develop large-scale or domain-specific knowledge bases or ontologies; development of domain or task focused tools, techniques, or methods; evaluations of AI tools, techniques or methods for domain suitability; unsuccessful attempts to apply particular tools, techniques or methods to specific domains (which shed insight on the applicability and limitations of the tool, technique or method); system architectures that work; scalability of techniques; integration of AI with other technologies; development methodologies; validation and verification; lessons learned; social and other technology transition issues.

The following questions will appear on the

Timetable for Authors

- ⇒ December 1, 2010 – January 31, 2011: Authors register on the IAAI web site
- ⇒ February 1, 2011: Electronic papers due
- ⇒ April 26, 2011: Camera-ready copy due at AAAI office

review form for emerging technology papers. Authors are advised to bear these questions in mind while writing their papers. Reviewers will look for papers that meet at least some (although not necessarily all) of the criteria in each category.

Significance: How important is the problem being addressed? Is it a difficult or simple problem? Is it central or peripheral to a category of applications? Is the tool or methodology presented generally applicable or domain specific? Does the tool or methodology offer the potential for new or more powerful applications of AI?

AI Technology: Does the paper identify AI research needed for a particular application or class of applications? Does the paper characterize the needs of application domains for solutions of particular AI problems? Does the paper evaluate the applicability of an AI tool or methodology for an application domain? Does the paper describe AI technology that could enable new or more powerful AI applications?

Innovation: Does the tool, technique, or method advance the state of the art or state of the practice of AI technology? Does the tool, technique, or method address a new or previously reported problem? If it is a previously reported problem, does the tool, technique, or method solve it in a different, new, more effective, or more efficient way? Does the reported work integrate AI with other AI or non-AI technologies in a new way? Does the work provide a new perspective on an application domain? Does the work apply AI to a new domain?

Content: Does the paper motivate the need for the tool or methodology? Does the paper adequately describe the task it performs or the problem it solves? Does it provide technical details about the design and implementation of the tool or methodology? Does the paper clearly identify the AI research results on which the tool or methodology depends? Does it relate the tool or methodology to the needs of application domains? Does it provide insights about the use of AI technology in general or for a particular application domain? Does it describe the development process and costs? Does it discuss estimated or measured benefits? Does it detail the evaluation method and results?

Evaluation: Has the tool or methodology been tested on real data? Has it been evaluated by end users? Has it been incorporated into a deployed application? Has it been compared to other competing tools or methods?

Technical Quality: Is the paper technically sound? Does it carefully evaluate the strengths and limitations of its contribution? Are the results described and evaluated? Are its claims backed up? Does it identify and describe relevant previous work?

Clarity: Is the paper clearly written? Is it organized logically? Are there sufficient figures and examples to illustrate the key points? Is the paper accessible to those outside the application domain? Is it accessible to those in other technical specialties?

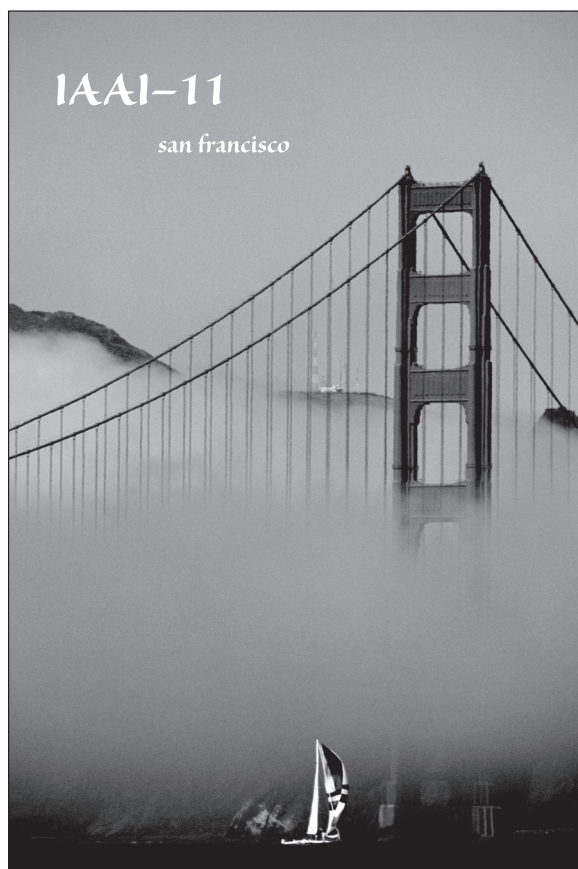
A copy of the review form for Emerging Application Case Study Papers will be posted on the IAAI website (www.aaai.org/Conferences/IAAI/iaai11.php) in the fall of 2010.

Submission Format

Electronic submissions are required. Papers must be in trouble-free, high resolution PDF format and formatted for United States Letter (8.5" x 11") paper. Submissions need to be in AAAI two-column format, (www.aaai.org/Publications/Author/author.php.) Deployed papers can be up to eight (8) pages. Emerging papers are limited to six (6) complimentary pages and two (2) optional additional pages at \$275 each.

Authors should register on the IAAI-11 web-based paper submission software (available December 1, 2010). A login and password, as well as detailed instructions about how to submit an electronic paper, will be sent to the author in a subsequent e-mail message. Authors must then submit a formatted electronic version of their paper through this software no later than Tuesday, February 1, 2011. We cannot accept papers submitted by e-mail or fax.

Submissions received after the deadline or that do not meet the length or formatting requirements detailed previously and at the IAAI-11 web site will not be accepted for review. Notification of receipt of the



electronic paper will be mailed to the first author (or designated author) soon after receipt. If there are problems with the electronic submission, AAAI will contact the primary author by e-mail. Papers will be reviewed by the Program Committee and notification of acceptance or rejection will be mailed to the contact author in late March. PDFs of accepted papers will be due on April 26, 2011. Authors will be required to transfer copyright.

Inquiries

Registration or clarification inquiries may be sent to AAAI at iaai11@aaai.org, 650-328-3123, or 650-321-4457 (fax).

- Daniel Shapiro (ISLE)
Conference Chair
- Markus Fromherz (Palo Alto Research Center)
Conference Cochair