

Reports of the AAAI 2012 Conference Workshops

Vikas Agrawal, Jorge Baier, Kostas Bekris, Yiling Chen, Artur S. d'Avila Garcez, Pascal Hitzler, Patrik Haslum, Dietmar Jannach, Edith Law, Freddy Lecue, Luis C. Lamb, Cynthia Matuszek, Hector Palacios, Biplav Srivastava, Lokendra Shastri, Nathan Sturtevant, Roni Stern, Stefanie Tellex, Stavros Vassos

■ The AAAI-12 Workshop program was held Sunday and Monday, July 22–23, 2012, at the Sheraton Centre Toronto Hotel in Toronto, Ontario, Canada. The AAAI-12 workshop program included nine workshops covering a wide range of topics in artificial intelligence. The titles of the workshops were *Activity Context Representation: Techniques and Languages*, *Cognitive Robotics*, *Grounding Language for Physical Systems*, *Human Computation*, *Intelligent Techniques for Web Personalization and Recommendation Systems*, *Multiagent Pathfinding*, *Neural-Symbolic Learning and Reasoning*, *Problem Solving Using Classical Planners*, and *Semantic Cities*. This article presents short summaries of those events.

Activity Context Representation: Techniques and Languages

Context-aware cognitive support requires activity and context information to be captured and, ever more often, moved across devices — securely, efficiently, and with multidevice interoperability. Context representation has received a lot of attention in the information technology community and in the industry among mobile vendors. 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.

The Activity Context Representation workshop was introduced by cochairs and coanchors Simon Scerri (Digital Enterprise Research Institute, National University of Ireland) and Wolfgang Woerndl (Technische Universität München, Germany) highlighting how activity-context-aware systems are critical to creating the next generation of digital workspaces, where activity context representation techniques are of central importance.

Gail Murphy (Tasktop and University of British Columbia) delivered a keynote on a powerful implementation of integrative tools for different knowledge workers (especially software developers), synchronizing across the enterprise context, enabling context capture, providing cognitive support for context switching, maintaining interaction histories with degrees of interest, and enabling sharing context of tasks.

Mills Davis (Project10x) delivered a keynote on Concept Computing, describing a paradigm shift from information-centered to knowledge-centered patterns of computing, highlighting the urgent business need for an architecture of user-centered software systems that are semantic model driven and activity context aware, where meaning is derived from a network of relationships. His keynote was followed by an open discussion, anchored by Boris Moltchanov (Telecom Italia) and Davis, on overcoming barriers to adoption and to discern levels of interest in different industry segments. Davis identi-

fied significant business interest in this area from defense, intelligence agencies, health, and finance. A future workshop on activity-context-aware systems architecture is being planned to address this need.

Ashish Kapoor (Microsoft Research) presented a keynote talk on context-aware face recognition by modeling context, considering cooccurrence, interconnection, events, and locations, using an inference solution that tries to probabilistically merge results from a classifier with a model of constraints. Boris Moltchanov (Telecom Italia) highlighted approaches to exploit mobile context-awareness-based services including social e-book reading and augmented reality on smartphones, using a context broker and context provider on a cloud-enabled service. Simon Scerri described an interoperable context representation called DCON that will be extended in the future to learn rules from user behavior. Wolfgang Woerndl described research results on how to infer user context from GPS logs, and proactively push recommendations to the user, with a context model that understands categories of context and a user activity recognition system that generates attributes for the context model. Rim Helaoui (University of Mannheim) presented an approach to recognizing activity at four levels (gestures, action, activity, situation) using probabilistic description logics for a smart house with sensors, using commonsense knowledge to resolve ambiguity, using a hybrid data-driven and knowledge-driven approach. Kai Kunze (University of Passau) highlighted an activity recognition model allowing dynamically configurable context-recognition systems. Kai highlighted the need for more open data sets for context recognition. Klaus Herrmann (University of Stuttgart) described a mechanism for collecting large-scale data from a public sensing system in an energy-efficient and user-friendly manner. Jeff Lockhart and Gary Weiss from Fordham University showed the impact of personalization on smartphone-based activity recognition and concluded that individual user data is critical to the personalization effort. Bostjan Kaluza (Jozef Stefan Institute, Slovenia) pro-

posed an ambient intelligence repository for activity recognition. Joey Chiao-yin Hsiao (National Taiwan University) showed how his team recognizes social engagement levels in the context of dyadic conversations. Naveen Nair (Indian Institute of Technology, Mumbai) presented challenges in learning optimal models for complex, first-order activity recognition. Genoveva Galarza Heredero (Infosys) presented a detailed manifesto and architecture for activity-context-aware systems that showed the features and detailed functions of each component. She showed a demonstration of the concept of this architecture in a patient care implementation.

Based on concrete outcomes from this workshop, we will start defining and exploring activity-context-aware system architectures in future workshops.

Lokendra Shastri (Infosys Limited, Workshop General Chair), Tim Finin (University of Maryland, Baltimore County), Henry Kautz (Rochester), Bo Begole (Samsung UX), and Munindar P. Singh (North Carolina State University) organized this workshop together with Vikas Agrawal. This report was written by Lokendra Shastri and Vikas Agrawal. The papers from the workshop were published as AAAI Technical Report WS-12-05.

Cognitive Robotics

The use of both software robots (softbots) and robotic artifacts in everyday life is on the upswing and we are seeing increasingly more examples of their use in society with commercial products around the corner and some already on the market. As interaction with humans increases, so does the demand for sophisticated robotic capabilities associated with deliberation and high-level cognitive functions. Combining results from the traditional robotics discipline with those from AI and cognitive science has and will continue to be central to research in cognitive robotics. As improvements in technology lead to more sophisticated sensors, actuators, and so on, together with a decrease in size, power consumption, and cost, the need for higher-level cognitive functions on these

more sophisticated robots is becoming more and more pressing.

Research in robotics has traditionally emphasized low-level sensing and control tasks including sensory processing, path planning, and manipulator design and control. In contrast, research in cognitive robotics is concerned with endowing robots and software agents with higher-level cognitive functions that enable them to reason, act, and perceive in changing, incompletely known, and unpredictable environments. Such robots must, for example, be able to reason about goals, actions, when to perceive and what to look for, the cognitive states of other agents, time, collaborative task execution, and so on. In short, cognitive robotics is concerned with integrating reasoning, perception, and action within a uniform theoretical and implementation framework.

This two-day workshop brought together researchers from a variety of subfields of AI in addition to researchers in the field of cognitive science. The 15 papers that were presented in the workshop focused on a broad range of topics from theoretical approaches and methodologies for various aspects of cognitive robotics to practical and challenging applications of cognitive robots in the real world.

Some of the topics discussed by the presenters include practical languages for reasoning about action and change, agent programming languages for high-level behavior, plan recognition, modeling visuospatial abilities, as well as the framework of conceptual spaces, the Soar cognitive architecture, and emotional intelligence. As far as practical applications of cognitive robotics is concerned, the discussions included an application in a personalized guided tour using cognitive robots (Personalized Guided Tour by Multiple Robots Through Semantic Profile Definition and Dynamic Redistribution of Participants by Anna Hristoskova, Carlos Aguero, Manuela Veloso, and Filip De Turck), a framework for modeling the skills of manufacturing robots (A Taxonomic Framework for Task Modeling and Knowledge Transfer in Manufacturing Robotics by Jacob Huckaby and Henrik I. Christensen), as well as a prototype robotic bartender (What Would

You Like to Drink? Recognising and Planning with Social States in a Robot Bartender Domain by Ron Petrick and Mary Ellen Foster).

In addition to the contributed presentations, the workshop included two invited talks from renowned members of the AI community that took place at the beginning of each day of the workshop. Michael Thielscher (University of New South Wales, Australia) talked about the research topic of general game playing in the context of cognitive robotics, and Peter Stone (University of Texas, USA) focused on the RoboCup competition, discussing its evolution in recent years and the major challenges that have been identified.

The discussions that were initiated by the invited talks and the paper presentations also continued in a panel session that took place near the end of the second day of the workshop. The panel included four panelists, Giuseppe De Giacomo (Sapienza University of Rome, Italy), Unmesh Kurup (Carnegie Mellon University, USA), Gerhard Lakemeyer (Aachen University of Technology, Germany), and Peter Stone (University of Texas, USA), and the topic of the discussion was practical challenges for cognitive robotics.

Three major themes emerged from a lively discussion among the panelists and the participants of the workshop identifying the following as major challenges: (1) the ability to form ad hoc teams of cognitive robots for collaboratively addressing a given situation, (2) the use of software worlds, such as video games, not only as a test bed for experimenting and evaluating approaches for cognitive robotics, but also as a self-standing application domain that becomes ubiquitous in the modern age of information, and (3) an effective account of perception that incorporates both symbolic and non-symbolic reasoning also taking into account a fusion of the available vast amount of information in existing knowledge bases.

Wolfram Burgard (Institute for Informatik, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany), Kurt Konolige (Industrial Perception, Inc., Palo Alto, USA), Maurice Pagnucco (University of New South Wales, Sydney, Australia), and Stavros Vassos

(National and Kapodistrian University of Athens, Athens, Greece) served as co-chairs for this workshop. This report was written by Stavros Vassos. The papers of the workshop were published as AAAI Press Technical Report WS-12-06.

Grounding Language for Physical Systems

Natural language is a powerful and intuitive modality for enabling humans to interact with physical systems such as robots. Understanding language about the external world requires the ability to extract a semantically meaningful representation from the language and map that representation to external, nonlinguistic percepts. This problem, referred to as the language grounding problem, has received substantial attention recently; in part because advances in robotics and sensing technology have enabled more robust sensing, manipulation, and simulation of the physical world, and in part as a result of recent advances in natural language-processing and formal representation systems.

The Grounding Language for Physical Systems workshop was a day of invited talks, a poster session, and discussion, bringing together researchers from a variety of fields, including computational linguistics, robotics, and planning. Invited speakers were selected to highlight current research on language grounding in each of these communities.

The workshop was opened by Stefanie Tellex presenting for Seth Teller about affordances as a representation for sharing mental models. In natural language learning processing and learning, Luke Zettlemoyer spoke about learning to understand dialog, and Ray Mooney presented work on following natural language route directions by learning a mapping between words and actions. Other speakers described recent work guiding robotic learning and understanding: Cynthia Matuszek described work on learning perceptual features and word meanings from a corpus of language paired with images, Matthias Scheutz of Tufts gave an overview of his work on dialogue understanding, and Stefanie Tellex pre-

sented her team's latest work on asking targeted questions using a metric based on entropy. Henry Kautz presented on a multimodal data set being developed at Rochester to integrate language and action.

The workshop included a poster session (abstracts for the posters appear in AAAI Press Technical Report WS-12-07). This provided an opportunity for more informal interactions, incorporating papers on learning to interpret natural language, learning from language, and the use of sensors and visual context either to learn or to interpret language. Algorithms and approaches were presented for grounding word meanings in perceptual features, task planning based on language cues, goal-based representation for actions, and the use of cognitive and continuous architectures to support language learning. The submitted papers provided an exciting showcase of upcoming work in this area, and several themes emerged. Using natural language to instruct robots or agents — using sensor or environment data to ground commands or to understand features of the world — was a common task for evaluating and utilizing language grounding. A second theme was language grounding for improved human-robot interactions, with several papers describing the use of grounded language to support flexible and socially appropriate interactions.

The workshop concluded with a lively discussion. The consensus was that while the field is not mature enough for a single shared task, more pairwise sharing of data sets would be beneficial, and several teams made plans to do this. Participants felt that bringing together members of different communities was informative and beneficial, and agreed that they would like to attend future workshops with the same focus as this one.

Cynthia Matuszek, Stefanie Tellex, Dieter Fox, and Luke Zettlemoyer served as organizers of the workshop. This report was authored by Stefanie Tellex and Cynthia Matuszek. The papers of the workshop were published as AAAI Press Technical Report WS-12-07.

Human Computation

Human computation is a relatively new research area that studies how to build intelligent systems that involve human computers, with each of them performing computation (for example, image classification, translation, and protein folding) that leverages human intelligence but challenges even the most sophisticated AI algorithms that exist today. With the immense growth of the web, human computation systems can now leverage the abilities of an unprecedented number of Internet users to perform complex computation. Various genres of human computation applications are available today, including games with a purpose (for example, the ESP Game) that generate useful data through gameplay, crowdsourcing marketplaces (for example, Amazon Mechanical Turk) that coordinate workers to perform tasks for monetary rewards, and identity verification systems (for example reCAPTCHA) that generate useful data through users performing computation for access to online content.

Over the past few years, we have observed a proliferation of related workshops, new courses, and tutorials, scattered across many conferences. The goal of this workshop is to address this fragmentation by bringing together academic and industry researchers from diverse subfields — machine learning, mechanism and market design, information retrieval, decision-theoretic planning, optimization, human-computer interaction — in a stimulating discussion of recent solutions to the core research questions in human computation and the future directions of this relatively new research area.

The 30 accepted papers and posters fall into three broad themes. The first theme highlights the design of human computation systems, studying how to use different motivations, workflows, and social norms to improve the quality of human computation. The second theme focuses on challenges of integrating human and machine intelligence in computation. Papers in this theme make progress toward achieving better integrated human-machine systems by leveraging their complementary strengths. The two invited talks further highlight these challenges and

progresses. Adam Kalai, in his talk “Programming by Example Revisited,” discusses the importance of breaking a task into small steps, each of which when carried out by humans is easy to understand by computers. Jeffrey Bigham, in his talk *Crowd Agents: Interactive Crow-Powered Systems in the Real World*, illustrates a new approach of enabling a dynamic crowd by providing the computational support needed to act as a single, high-quality agent. Finally, there was a wide array of new tools and platforms that address the practical challenges faced by human computation systems in the real world.

This workshop marks the fourth year that we have organized a workshop on human computation. The number of submissions, as well as their diversity, has grown steadily over the past four years. In 2013, the workshop will transit into the first AAAI Conference on Human Computation. The workshop ended with a business meeting led by Eric Horvitz and engaged with a large, enthusiastic audience to discuss opportunities and practical issues surrounding the new conference.

Yiling Chen (Harvard) chaired the workshop, along with Luis von Ahn (CMU), Panos Ipeirotis (NYU), Edith Law (CMU), and Haoqi Zhang (Harvard) serving as organizers. This report was written by Edith Law and Yiling Chen. Papers of this workshop were published as AAAI Press Technical Report WS-12-08.

Intelligent Techniques for Web Personalization and Recommendation Systems

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 partic-

ular 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 researchers and practitioners and researchers from the partially overlapping fields of web personalization and recommender systems to discuss current and emerging topics in the respective fields and to foster an exchange of ideas and experiences. The workshop was opened with an invited talk by Blaz Fortuna from the Jozef Stefan Institute at Ljubljana, Slovenia. In his talk, he reported on the design and implementation of a real-world news recommendation system, which considers various types of information in the recommendation process and is capable of handling up to a few hundred page views per second.

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 ranging from the exploitation of social web data (for example, from microblogging platforms) for personalization over recent topics such as group recommendation and location-based advertising to the automatic classification of web pages for personalization.

The discussions after the technical paper presentations were focused on questions related to the exploitation and integration of different sources of knowledge in the personalization and recommendation process as well as on

methodological questions concerning experimental designs for the evaluation of systems and algorithms including the role of user studies and real-world live evaluations.

The workshop was organized and cochaired by Dietmar Jannach (TU Dortmund, Germany), Sarabjot Singh Anand (University of Warwick, UK), Bamshad Mobasher (DePaul University, USA) and Alfred Kobsa (University of California, Irvine, USA). Dietmar Jannach authored this report. The papers of the workshop were published as AAAI Press Technical Report WS-12-09.

Multiagent Pathfinding

Recently, there has been a growing interest in multiagent pathfinding (MAPF). The problem is to compute a path for each agent from an initial to a goal location without conflicting with other agents, often aiming to minimize a cost function, such as elapsed time or throughput. Applications include vehicle fleet coordination, computer games, robotics, and various military scenarios. Research on this problem covers the full spectrum between purely theoretical and fully application-driven work. Thus, related papers have appeared in multiple venues, including AIJ, JAIR, AAAI, IJCAI, ICRA, IROS, ICAPS, and SoCS. Theoretical work is also found outside the AI and robotics communities. Consequently, similar concepts were developed in different subcommunities, using varying terminology.

The main goals of the workshop were to bring together researchers working on multiagent path planning from different communities, and encourage collaboration between them; familiarize researchers from different areas with the varying contributions on this problem; standardize terminology and develop a taxonomy for different variants; and present the state of the art and discuss open challenges.

The first half of the workshop featured mostly papers that address a discrete variant of the multiagent pathfinding problem, where there is a finite number of locations where the agents can be located. Presentations by Trevor Standley (Google) and Guni Sharon (Ben-Gurion University) covered state-of-the-art solutions utilizing

search-based approaches. Malte Helmert's presentation (joint work with Gabriele Roger, Freiburg) highlighted a theoretical work by Kornhauser and colleagues in 1984, which can be used to construct a complete (but suboptimal) multiagent pathfinding algorithm in polynomial time. Yingjin Yu (UIUC) presented how ideas from network flow can be used to solve multiagent pathfinding problems with desired properties.

The second half of the workshop primarily featured talks that address the continuous variant of the problem. This started with a discussion by Kostas Bekris (Rutgers) on how discrete algorithms could be potentially used to solve continuous instances and included a presentation by Dinesh Manocha (UNC Chapel Hill) on reciprocal collision avoidance. This portion of the workshop also described application-specific work, including role assignment and its use in robot soccer (by Patrick Macalpine, University of Texas, Austin), as well as work on planning the motion of multiple quadrotors through reciprocal collision avoidance (by Daman Bareiss, University of Utah).

The workshop included a set of poster presentations, which varied from optimal search-based solutions to control-based methodologies.

At the conclusion of the workshop, the participants discussed how best to encourage and broaden the collaboration of researchers working on this problem. This includes the possibility of pursuing the colocation of a future related event with different conferences to reach the different communities interested in this problem.

Roni Stern, Kostas Bekris, Nathan Sturtevant, and Ariel Felner served as chairs of this workshop, and the first three authored this report. The papers of the workshop were published as AAAI Press Technical Report WS-12-10.

Neural-Symbolic Learning and Reasoning

The Neural-Symbolic Learning and Reasoning workshop at AAAI-12 was the latest edition of a series started at IJCAI-05 and then organized yearly at ECAI-06, IJCAI-07, ECAI-08, IJCAI-09, AAAI-10, and IJCAI-11. The AAAI workshop

provided a venue for active researchers in the areas of machine learning, neural computation, knowledge representation, cognitive science, logic, and artificial intelligence to share experiences.

This year was particularly relevant for artificial intelligence and computer science. Several celebrations, exhibitions, and symposia have celebrated Alan Turing's 100th birthday. Turing not only defined the foundational basis of computer science but also created the fields of artificial intelligence, machine learning, and neural computation. After more than 50 years of intense research, artificial intelligence still faces huge challenges in its quest to develop truly intelligent systems. Such systems are required to learn and adapt to changes in the environment in which they operate and to reason about common-sense knowledge in ways that can control the accumulation of errors. Recent developments in the area of neural-symbolic computation and machine learning 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 started with a keynote talk by Jude Shavlik on 25 years of combining symbolic and numeric learning. Shavlik addressed the achievements and challenges machine learning faced over a quarter of a century, outlining

developments started on explanation-based learning and knowledge-based neural through knowledge-based support vector machines and Markov logic networks.

The talks spanned from the foundations of symbolic reasoning and machine learning to applications in temporal reasoning, software engineering, speech recognition, visual intelligence, vehicle control, and cognitive science. Alan Perotti (jointly with Guido Boella, Artur d'Avila Garcez, and Daniele Rispoli) presented the paper Neural-Symbolic Rule-Based Monitoring. Perotti presented a neural-symbolic system for monitoring traces of observations in software systems. He showed how a rule-based neural network system (RNNS) can perform trace monitoring effectively and analyzed the system's performance.

Jim Prentzas (jointly with Ioannis Hatzilygeroudis) presented the paper Efficiently Merging Symbolic Rules into Integrated Rules. They explained how neurules (a type of neural-symbolic rules that integrate neurocomputing and production rules) exhibit characteristics such as modularity, naturalness, and ability to perform interactive and integrated inferences.

Christian Huyck presented the paper Vowel Recognition in Simulated Neurons, in which he described a neural simulation that learns to categorize three vowel sounds. This novel model is probably the best neuropsychological model of vowel recognition, because it is the only one that uses simulated neurons starting with reasonable neural input.

Tsvi Achler presented the paper Towards Bridging the Gap between Pattern Recognition and Symbolic Representation within Neural Networks. Achler argued that underlying symbolic representations are opaque within neural networks that perform pattern recognition. He showed that by implementing network dynamics differently, during the testing phase instead of the training phase, pattern recognition can be performed using symbolically relevant weights. Achler argued that this is an important step toward the merging of neural-symbolic representation, memory, and reasoning with pattern recognition.

Leo de Penning (joint work with Richard Den Hollander, Henri Bouma, Gertjan Burghouts, and Artur d'Avila Garcez) presented A Neural-Symbolic Cognitive Agent with a Mind's Eye. Motivated by the DARPA Mind's Eye program, de Penning described a neural-symbolic cognitive agent that integrates neural learning, symbolic knowledge representation, and temporal reasoning in a visual intelligent system that can reason about actions of entities observed in video. Their results illustrate that the system is able to learn and represent the underlying semantics of the actions from observation and use this for several visual intelligent tasks, such as recognition, description, anomaly detection, and gap filling.

Yury Vinokurov (joint work with Dean Wyatte, Christian Lebiere, Randall O'Reilly, and Seth Herd) presented the paper Unsupervised Learning in Hybrid Cognitive Architectures. Vinokurov presented a new model of unsupervised learning in the hybrid SAL (Synthesis of ACT-R and Leabra) architecture. The presented model follows the hypothesis that higher evaluative cognitive mechanisms can serve to provide training signals for perceptual learning. This model also provides a basis for implementing unsupervised learning and symbol grounding in dynamic circumstances by agents navigating through and operating in such environments.

Dragos Margineantu's (joint work with Tomas Singliar) talk was named Scalable Inverse Reinforcement Learning via Instructed Feature Construction. Margineantu presented a procedure to scale up inverse reinforcement learning (IRL) by eliciting good IRL basis functions from the domain expert. He also proposed a new paradigm for modeling limited rationality. Unlike traditional models of limited rationality that assume an agent making stochastic choices with the value function being treated as if it is known, Margineantu proposes that observed irrational behavior is actually due to uncertainty about the cost of future actions.

The workshop ended with a discussion session on the future of machine learning in Alan Turing's centenary. Participants highlighted the impor-

tance of Turing's work both as an artificial intelligence pioneer, but also his fundamental contributions to the establishment of machine learning and neural computation. Logic plays a central role in Turing's most fundamental works on computability and AI. We draw inspiration from his outstanding work, and we believe that symbolic (logical) knowledge and neural and statistical learning are fundamental to the development of truly intelligent systems. In the tradition of the NeSy workshop series, a number of suggestions and challenges for further research have been discussed. In particular, applications of neural-symbolic computation in visual intelligence, software engineering, temporal learning, and language processing have recently shown that the methodology is highly effective in complex domains and application areas.

Artur d'Avila Garcez, Pascal Hitzler, and Luis Lamb were the workshop cochairs and authored this report. The papers of the workshop were published as AAAI Press Technical Report WS-12-11.

Problem Solving Using Classical Planners

The performance of classical planners has advanced greatly over the last two decades, and many now try to exploit this, through compilations and other methods of reuse, to solve a much wider range of problems. The workshop aimed to showcase inspiring new uses of classical planners and to bring together planning researchers with those who try to use planners to solve new problems.

Papers presented at the workshop covered many different angles. The problems they addressed ranged from planning under uncertainty (contingent, probabilistic, or conformant), to applications in service composition, system reconfiguration and cyber security, to rather unusual applications such as plan critiquing and even generating new planning problems. An invited talk, by Ken Pickering (Core Security), detailed how classical planning is one key ingredient in computer network security testing. Approaches presented ranged from pure compilation (trans-

lating a problem into PDDL to solve it with a classical planner) to embedding classical planning techniques inside dedicated algorithms.

Yet, we were interested to note a few recurrent themes, particularly in application papers. One was the desire to generate sets of plans that exhibit diversity. Although some work on plan diversity exists, this still seems to be a relevant, and nontrivial, challenge for planning research. Even how to define what is interesting diversity in a set of plans seems to be an open question.

Another was the difficulty of formulating problems in such a way that planners can effectively deal with them. While nearly all planners now use PDDL for input, there is often a large gap between the full set of features available in PDDL and that supported by planners. It was also noted that the semantics of PDDL, because of the evolution of the language over time, are not well documented, and sometimes ambiguous. This leads to planners taking different interpretations of (particularly the less frequently used parts of) PDDL, causing difficulties for users. Clarifying and documenting PDDL is an important project for the whole planning community.

Planner users prefer a rich modeling language, while planning research often focuses on the propositional STRIPS setting. Even though many non-STRIPS features of PDDL can be removed through effective reformulations, these encoding tricks are often not obvious, and knowledge of them is not widespread. Such formulations are also often seen as “unnatural,” which can be a serious impediment. In some applications, domain models are very large, built by teams of domain experts, and remain in use over a long time. Thus, aspects such as modularity and maintainability of domain models are important for planning to be viable. Some features, most notably conditional effects, cannot always be compiled away efficiently, and planner users clearly desire better support for them. Extending current planning methods to properly handle conditional effects is nontrivial and presents a challenge for planning research. Grounding (as done by most current planners) was also identified as an obstacle to applicability

on large problems. What the right level of complexity at the “interface” to planners should be — that is, on which side different features should be handled, through compilation or other means — is likely to be an ongoing debate in the community for some time.

The workshop was organised by Hector Palacios, Jorge Baier, and Patrik Haslum who also authored this report. The papers are available as AAAI Technical Report WS-12-12.

Semantic Cities

In a semantic city, available resources are harnessed safely, sustainably, and efficiently to achieve positive, measurable economic and societal outcomes. Enabling city information as a utility, through (1) a robust, expressive, dynamic, scalable, and, critically, a sustainable technology and (2) a socially synergistic ecosystem, could drive significant benefits and opportunities. Data, and then derived information and further derived knowledge, from people, systems, and things, are the single most scalable resource available to city stakeholders to reach the objective of semantic cities.

Two major trends are supporting semantic cities — open data and the semantic web. Open data is the idea that data should be accessible for everyone to use and republish as they wish, without restrictions from copyright, patents, or other mechanisms of control. A number of cities and governments have made their data publicly available, prominent being London (UK), Chicago (USA), Washington DC (USA), and Dublin (Ireland). The semantic web as the technology to interconnect heterogeneous data has matured and it is being increasingly used in the form of linked open data and formal ontologies. Thus, a playfield for more AI research-driven technologies for cities has emerged for example, scalable, efficient, robust, optimal AI techniques.

The workshop had a rich program highlighting how to build semantic cities and provide societal benefits. It consisted of two invited talks, six regular papers, two short papers and demonstrations, and a panel attracting

speakers and attendees from academia, industry, and the open data movement. Here, we discussed (1) web technologies related to linked data and open data, (2) AI topics such as ontology design, description logic-based reasoning, stream reasoning, and optimization, (3) use cases related to semantic-augmented public services, citizen-driven applications, traffic regulation, the next generation of the U.S. information service 311, earthquake disaster exploration, smart grids, open government going beyond economic growth to tackling corruption, air traffic management, finance, energy, and telecommunications, (4) database topics such as data streams, large-scale data integration, and (5) vertical research challenges related to privacy and security.

Questions related to the level of city data representation, that is, expressivity and integration, were discussed at length, and the participants finally reached the conclusion that the trade-off between expressivity and tractability needs to be addressed case by case. Any generalization is then very difficult to be reached. As an example, the participants pointed out the needs for data streams to be semantically described in a lightweight way, as data streams require fast processing and reasoning.

The workshop participants also discussed how interest in semantics for cities and public bodies is growing while pointing out the lack of best practices to reach the level of semantic cities. In this respect the participants agreed that the workshop was useful and successful. They were most enthusiastic about the application-driven presentation of various contributions of the workshop although some participants asked for more applications, demonstrations, and views on business impacts, for example, by having city representatives for the next series. They also clearly agreed on the needs and benefits of applying and revisiting AI techniques for solving real-world city problems such as traffic regulation, traffic road congestion, heterogeneous data integration in a city context, and air traffic management. It was noted that new research challenges arise from these real-world domains. The workshop also had a multidisciplinary

International Conference on Intelligent User Interfaces



Santa Monica, CA, USA
March 19-22, 2013

Homepage: <http://iuiconf.org/>
Twitter: @IUIConference
Facebook: [facebook.com/IUI2013](https://www.facebook.com/IUI2013)

IUI 2013 is the annual meeting of the intelligent interfaces community and serves as the principal international forum for reporting outstanding research and development on intelligent user interfaces.

We invite original paper submissions that describe novel user interfaces, applications, interactive and intelligent technologies, empirical studies, or design techniques.

Topics	
Multi-modal input • Natural language • Gesture • Face recognition	
Ubiquitous computing • Smart environments	
Intelligent assistants • Multiuser environments	
Affective, social and aesthetic interfaces • Social networks and collaboration	
Personalization • Recommender systems • User modelling	
Plan recognition • Knowledge-based interface generation	
Agent-based paradigms • Example- and demonstration-based interfaces	
User studies for IUIs	
Visualization • Novel interfaces for linked data • Interfaces for large ontologies	

Important Due Dates	
Papers (abstract)	Oct 15, 2012
Papers (final)	Oct 22, 2012
Poster	Jan 9, 2013
Workshop	Oct 22, 2012
Student Consortium	Nov 16, 2012
Demo	Jan 9, 2013




21st International Conference on Case-Based Reasoning
July 7th – 10th 2013, Saratoga Springs, NY, USA

We invite you to participate in the 21st International Conference on Case-Based Reasoning, the premier international meeting on CBR. In addition to the technical program, ICCBR 2013 will host the ICCBR workshop series, industry day, Doctoral Consortium, and the 6th Computer Cooking Competition (CCC).

Program Chairs:

- Sarah-Jane Delany (Dublin Institute of Technology, Ireland)
- Santiago Ontañón (Drexel University, USA)

Local Chair:

- William E. Cheetham (General Electric, USA)

Industry Day Chairs:

- William E. Cheetham (General Electric, USA)
- Frode Sørmo (Verdande Technology, Norway)

Sponsorship Chair:

- David Aha (Naval Research Laboratory, USA)

Workshop Chairs:

- Michael Floyd (Carleton University, Canada)
- Jonathan Rubin (Palo Alto Research Center, USA)

Doctoral Consortium Chair:

- Thomas Roth-Berghofer (University of West London, UK)
- Rosina Weber (Drexel University, USA)

Conference website: <http://www.iccbr.org/iccbr13>

In cooperation with AAAI

dimension, for example, the fusion of database and AI techniques for tackling problems city managers are facing every day.

Participants agreed that they would like the workshop to be held again at AAAI as it fit perfectly the multidisciplinary feature required to tackle semantic cities challenges. They also shared the idea of having a tutorial to improve the external recognition of semantic cities in an AI context.

Biplav Srivastava, Freddy Lecue, and Anupam Joshi served as coauthors of this workshop. Biplav Srivastava and Freddy Lecue were the authors of this report. The papers of the workshop were published as AAAI Press Technical Report WS-12-13.

Vikas Agrawal is a research analyst at the Center for Knowledge Driven Intelligent Systems at Enterprise Technology Research Labs at Infosys Limited.

Jorge Baier is an assistant professor at the Pontificia Universidad Católica de Chile.

Kostas Bekris is an assistant professor at Rutgers University, New Jersey.

Yiling Chen is an associate professor of computer science at Harvard University.

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

Pascal Hitzler is an associate professor of computer science at Wright State University, Dayton, OH.

Patrik Haslum is a research fellow at the Australian National University.

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

Edith Law is a Ph.D. candidate in the Machine Learning Department at Carnegie Mellon University.

Freddy Lecue is a research scientist at IBM Research – Smarter Cities Technology Centre, Dublin, Ireland.

Luis C. Lamb is a professor of computer science at the Federal University of Rio Grande do Sul, Porto Alegre, Brazil.

Cynthia Matuszek is a Ph.D. candidate at the Computer Science and Engineering Department, University of Washington.

Hector Palacios holds a postdoctoral position at the Universidad Carlos III de Madrid.

Biplav Srivastava is a senior research scientist at IBM Research – India.

Lokendra Shastri heads the Enterprise Technology Research Labs and the Center for Knowledge Driven Intelligent Systems at Infosys Limited.

Nathan Sturtevant is an assistant professor at the University of Denver, USA.

Roni Stern is a postdoctoral fellow at Ben Gurion University of the Negev, Israel.

Stefanie Tellex is a research scientist at the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

Stavros Vassos is a postdoctoral research associate at the Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece.



**23rd International Conference on
Automated Planning and Scheduling**
Rome, Italy 10-14 June 2013



The 23rd International Conference on Automated Planning and Scheduling

ICAPS is the premier forum for exchanging news and research results on theory and applications of intelligent planning and scheduling technology.

- Paper submissions to workshops is open (see workshops list on the web)
- Co-located with the ICAPS summer school on Automated Planning & Scheduling

We invite you to participate at
ICAPS 2013

<http://icaps13.icaps-conference.org/>



ICWSM

The Seventh International AAAI Conference on Weblogs and Social Media (ICWSM-13)

MIT Media Lab
Cambridge, Massachusetts, USA
July 8–11, 2013

The broad goal of ICWSM is to increase understanding of social media in all its incarnations. Submissions describing research that blends social science and computational approaches are especially encouraged.

Paper Submissions Due February 6, 2013

www.icwsml.org/2013

CAMBRIDGE

Hot New Titles in AI from Cambridge University Press


E-books
Available
for most
titles!

Networked Life

20 Questions and Answers

Mung Chiang

\$45.00: Hb: 978-1-107-02494-6: 488 pp.



The Cambridge Handbook of Cognitive Science

*Edited by Keith Frankish and
William Ramsey*

\$90.00: Hb: 978-0-521-87141-9: 352 pp.

\$29.99: Pb: 978-0-521-69190-1




Thinking through the Body

Essays in Somaesthetics

Richard Shusterman

\$99.00: Hb: 978-1-107-01906-5: 374 pp.

\$32.99: Pb: 978-1-107-69850-5



Data-Intensive Computing

Architectures, Algorithms, and Applications

*Edited by Ian Gorton and
Deborah K. Gracio*

\$70.00: Hb: 978-0-521-19195-1: 312 pp.




Machine Learning

The Art and Science of Algorithms that
Make Sense of Data

Peter Flach

\$120.00: Hb: 978-1-107-09639-4: 409 pp.


\$60.00: Pb: 978-1-107-42222-3



Causality, Probability, and Time

Samantha Kleinberg

\$99.00: Hb: 978-1-107-02648-3: 272 pp.



Computer Vision for Visual Effects

Richard J. Radke


\$70.00: Hb: 978-0-521-76687-6: 408 pp.



Plasticity in Sensory Systems

*Edited by Jennifer K. E. Steeves and
Laurence R. Harris*

\$120.00: Hb: 978-1-107-02262-1: 304 pp.



Prices subject to change.



**CAMBRIDGE
UNIVERSITY PRESS**



www.cambridge.org/us/computerscience

800.872.7423

@cambUP_maths