

Reports on the 2006 AAAI Fall Symposia

*Joshua Bongard, Derek Brock, Samuel G. Collins,
Ramani Duraiswami, Tim Finin, Ian Harrison,
Vasant Honavar, Gregory S. Hornby, Ari Jonsson,
Mike Kassoff, David Kortenkamp, Sanjeev Kumar,
Ken Murray, Alexander I. Rudnický,
and Goran Trajkovski*

- The American Association for Artificial Intelligence was pleased to present the AAAI 2006 Fall Symposium Series, held Friday through Sunday, October 13–15, at the Hyatt Regency Crystal City in Washington, DC. Seven symposia were held. The titles were (1) Aurally Informed Performance: Integrating Machine Listening and Auditory Presentation in Robotic Systems; (2) Capturing and Using Patterns for Evidence Detection; (3) Developmental Systems; (4) Integrating Reasoning into Everyday Applications; (5) Interaction and Emergent Phenomena in Societies of Agents; (6) Semantic Web for Collaborative Knowledge Acquisition; and (7) Spacecraft Autonomy: Using AI to Expand Human Space Exploration.

Aurally Informed Performance: Integrating Machine Listening and Auditory Presentation in Robotic Systems

This symposium brought together a number of researchers who are concerned with performance issues that robots face that depend, in some way, on sound. Many commercially marketed robotic platforms, as well as others that are moving from the laboratory into specialized public settings, already have rudimentary speech communication interfaces, and some are even being engineered for specific types of auditory tasks. In general, though, the ability of robots to monitor the auditory scene before them and to execute interactive behaviors informed by the interpretation or production of sound information remains far behind the broad and mostly transparent skills of human beings. It is an easy thing, for instance, for people to discern on the basis of audition alone who a familiar voice is, where the voice is located in the environment, and to act on other aspects of the au-

ditory scene before them, such as noise, informative sounds, or the need for proximity or loudness to facilitate verbal communication. When these auditory skills are integrated with people's other perceptual and reasoning abilities, substantial capacities for performance and interaction arise.

The design goals for robot audition and utterance of information by sound, though, are not just those that correspond to human skills. Machine auditory sensing can be designed, in certain ways, to be more capable and acute than human hearing, and going beyond speech, robotic auditory displays can be engineered to render nonspeech auditory information in and for a variety of manners and purposes. Thus, a substantial interaction design space arises when different modes of human-robot interaction are augmented by conventional and enhanced auditory functions.

Since the idea of "aurally informed performance" can be thought of as a two-sided proposition, involving both listening and presentation behaviors, the symposium was organized to focus on these themes separately and then conclude with a session on integrated systems. Much of the contributed research fit easily into this division. On each of the two main days, we began with an outline of recent research trends in the day's topic, "Listening" on the first day and "Presentation" on the second, and then heard contributed talks from participants. Afternoons were devoted to critical discussions of talks given in the morning and examination of a relevant philosophical question about the nature and role of sound as information in the design context of robotics.

The outline of trends in machine listening covered developments in sound localization and techniques for recognizing and extracting information from sound. Papers given under this day's theme made several contributions in these areas. Novel methods for classifying acoustical environments, localizing sounds with head-related transfer functions, organizing sounds with similar meanings, and perceiving and synthesizing speech were presented, and a chip being engineered for real-time classification of

affect in speech was described. In the afternoon we debated whether or not auditory features are universal with regard to the attribution of meaning. In addition, the idea of auditory primitives was considered as a conceptual computational analogue of speech primitives.

Our daylong session on the theme of aurally informed presentation began with an examination of this notion as referring not only to the production of sound information but also to the presentation of actions in response to auditory perceptions. Our outline of research trends in this area covered recent developments in speech and nonspeech auditory displays and nonauditory actions informed by sound. The day's talks ranged from the design of robotic utterances of nonspeech auditory information for affective and collaborative purposes, meeting users' needs as listeners, and continuous auditory feedback for user control in teleoperation to the analysis of musical features for coordinating robotic dance movements. In the afternoon we debated the nature of presentation in auditory interfaces, particularly when sound information must be meaningful in isolation, and the importance of perceptual ecologies users are likely to be familiar with and rely on in multimodal contexts involving sound.

The final day opened with a brief look at research involving audition and speech in developmental robotics and concluded with two talks on recent work with integrated robotic auditory interaction systems. The first described the integration of real-time auditory processing hardware capable of localizing, separating, and tracking several simultaneous speech sources with a dialogue manager on a mobile robot. The other described work on the problem of recognizing which agent is being addressed in a spoken dialogue system involving multiple human and robotic speakers and hearers being developed for future planetary ground exploration missions.

Derek Brock, Ramani Duraiswami, and Alexander I. Rudnicky cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-01 from AAAI Press.

Capturing and Using Patterns for Evidence Detection

Pattern-based analysis of data plays an increasing role in several important domains. In homeland security applications, it is used to detect evidence of threats and to predict threats before they completely mature. In social networks it is used to detect the relative importance of particular vertices or paths in the network. The goal of this symposium was to bring together researchers from diverse backgrounds to discuss the challenges they encounter in capturing and using patterns for evidence detection and the techniques they adopt to address those challenges. The symposium included 33 participants representing industry as well as academic and governmental institutions, and the 21 presentations covered a wide range of topics.

The presentations described diverse and innovative applications of pattern-based analysis. One application involves detecting and resolving name variations for data normalization. This ubiquitous problem arises when nicknames or spelling variations arise among the multiple references to a single entity in the data to be analyzed. A second application involves tracking relevant business news in support of a corporate radar. This very ambitious project attempts to detect patterns of interest in, for example, web-based news relevant to a given corporation, such as news about competitors, suppliers, and customers of the corporation. A third novel application involves a social-network analysis of the coaches of the National Football League. This entertaining domain was used to illustrate, among other things, the use of mentoring relationships among coaches to predict the success of a new head coach. Other interesting and innovative applications were described, including event recognition from airborne imagery and situation modeling during teamwork training using simulated environments.

Several papers focused on reasoning or representation issues. One theme common to several discussions involved determining the proper role for

abductive inference and how it and other nondeductive inference methods are best supported and integrated with deductive methods. A second theme involved the ubiquitous issue of knowledge representation: determining appropriate conceptualizations of patterns and domain models. One prominent thread of the symposium discussed a laboratory for evaluating link-analysis tools. The approach involves generating synthetic data so that ground truth is known but obscured from the tools and then scoring the performance of the link-analysis tools tasked with detecting target situations obscured in the synthetic data. One very interesting thread of the symposium discussed a mathematical model used to assess the complexity of the "connect the dots" (homeland security) problem. The model, based on a very intuitive jigsaw-puzzle metaphor, supports analyses on the solvability of the problem and the impact of various assumptions about the volume of data, the signal-to-noise ratio, and the number of independent analysts involved. While computer-human interface issues were not the focus of the symposium, they are essential in fielding link-analysis tools. Often the results of pattern matching can include a large amount of data. One important interface issue discussed involves how link-analysis results can be best presented to human users without overwhelming the user. Consideration of these and many other issues supported an enthusiastic discussion among the symposium participants.

Ken Murray and Ian Harrison cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-02 from AAAI Press.

Developmental Systems

During the process of biological development, structures as complex as the limbs, eyes, heart, and brain are constructed from millions of cells. Such complexities have captured the imaginations of scientists and philosophers for many years; but just how do these cells give rise to such structures? How is cell behavior controlled and by

what? How does an otherwise homogeneous structure give rise to such heterogeneity? These incredible feats of construction are testimony to the abilities of evolution by natural selection. Such questions capture the remarkable self-organization that occurs during biological development; they could also easily reflect much new research that is emerging under the banner of AI.

Artificial intelligence and developmental biology overlap on some quite important topics. One such topic is that of construction. Constructing robust complex adaptive systems in a self-organizing manner is a notoriously difficult problem. Inspired by biological development and engineering design, one promising approach to such problems is that of developmental systems: the use of computational models of development for constructing complex systems. In recent years, researchers have been investigating, with much success, methods of construction that overcome fundamental issues of adaptation, evolvability, scalability, modularity, self-organization, and self-repair. The developmental systems symposium focused on development and self-organizing principles that lead to the emergence of complex systems. New results were presented that explored mechanisms of self-repair and robust stochastic development of useful structures.

Sanjeev Kumar, Gregory S. Hornby, and Joshua Bongard cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-03 from AAAI Press.

Integrating Reasoning into Everyday Applications

Applications such as spreadsheets, search engines, and e-mail clients have become an integral part of modern everyday life. The user bases of these applications number in the hundreds of millions of users. This symposium was concerned with bringing the power of AI reasoning systems to aid such everyday applications.

This symposium was a sequel to the Workshop on Logical Spreadsheets, which was held at Stanford University in September 2005. Thus, several of

the presentations discussed logical spreadsheets, or spreadsheets in which the formula language is generalized from arithmetic formulae to logical formulae. Such spreadsheets excel in symbolic reasoning in addition to numerical reasoning. Bob Balzer, Iliano Cervesato, and David Van Brackle gave presentations on their respective systems.

Semantic e-mail was another hot topic at the symposium. Luke McDowell discussed semantic e-mail processes, which are declarative workflows for e-mail clients. Luke's system can automate tasks such as organizing a meeting with several participants and organizing a balanced potluck dinner, and it is currently at use at the University of Washington. Michael Kassoff discussed semantic e-mail addressing, or sending e-mails to people or groups of people based on their semantic features, rather than sending to their e-mail addresses.

Search engines were a final common topic at the symposium. Michael Witbrock of Cycorp discussed the natural language querying capability of Cyc, which uses lightweight parsing techniques coupled with a deep knowledge base of word senses to quickly determine candidates to answer a user query. Andrew Gordon discussed how logical inferences can supplement statistical text processing to achieve higher quality result sets to textual queries.

Michael Kassoff chaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-04 from AAAI Press.

Interaction and Emergent Phenomena in Societies of Agents

Multiagent systems cross disciplinary boundaries by focusing on society and cultures as emerging from interactions of autonomous agents. Poised at the intersections of AI, cybernetics, sociology, semiotics, and anthropology, this strand of multiagent systems research enables a powerful perspective illuminating not only how we live and learn but also, through focusing on emergence, how we anticipate the future. Whereas multiagent systems have

been extremely helpful in solving engineering problems, much of what we find exciting lies in their application to contemporary human life. In particular, the symposium focused on self-constituting systems and networks composed of human and non-human agents characteristic of emergent cybercultures, including e-commerce, e-learning, as well as other human-nonhuman agent systems in law, science, art, and online interactions of all kinds. We did not only share insights and experiments in multiagent systems composed of robot and software agents but also discussed at large the hybridity formed at the junction of the human and nonhuman.

At the 2006 AAAI Fall Symposium "Interaction and Emergent Phenomena in Societies of Agents," we invited researchers from all disciplines to dialogue, from their own viewpoint, on emergence. Some 25 AI researchers, psychologists, lawyers, economists, and artists from academia, government, and the industrial sector spent 2.5 days discussing a heavily packed agenda. The research presented both at the symposium and in the technical report is both cross-disciplinary and multidisciplinary and sheds light on different aspects of interaction and emergent phenomena in multiagent societies. We believe that we created new network assemblages of variegated agents of researchers and their techniques out of which may arise new perspectives on heretofore parochial questions in our respective disciplines.

All papers painted a comprehensive picture of the status of the art of this emergent discipline. Lazlo Gulyas (AITIA, Budapest, Hungary), Gyorgy Korniss (Rensselaer Polytechnic Institute, Troy, NY), and especially Mark Bickhard (Lehigh University, Bethlehem, PA) gave omnilateral views on the status of the study of emergence. The talk of Melanie Baljko and Nell Tenhaaf (York University, Toronto, Canada) titled "Different Experiences, Different Type of Emergence: A Life Sculpture Designer, Interactant, Observer" presented their interactive art sculpture that embodies theoretical concepts in interaction covered by the other speakers, and especially in the

talk of Erik Baumer (University of California-Irvine, Irvine, CA), on migrating virtual agents, and the human-computer interaction focused talk of Georgi Stojanov (American University in Paris, France).

About a fifth of the participants came from the Cognitive Agency and Robotics Laboratory at Towson University (Towson, MD), directed by Goran Trajkovski. Their talks on interactive (virtual and robotic) agents in homogeneous and heterogeneous environments gave interactivist perspectives on all topics covered, illustrated with ongoing experiments (POPSICLE, Izbushka, Baba-Yaga, and so on).

We feel that we established strong professional and personal connections with the whole group at the symposium. As a result, we are working together on a follow-up fall symposium in 2007 or 2008 and are in the process of finalizing a book proposal on interaction and emergence to be edited by the symposium organizers.

Goran P. Trajkovski and Samuel G. Collins cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-05 from AAAI Press.

Semantic Web for Collaborative Knowledge Acquisition

Recent advances in computing and communications, together with the rapid proliferation of ubiquitous networked information sources, sensors, and services, present unprecedented opportunities and challenges in integrative and collaborative analysis and interpretation of distributed, autonomous (and hence, inevitably semantically heterogeneous) information in virtually every area of human activity. Some compelling examples of such opportunities are found in emerging data-rich disciplines such as bioinformatics, environmental informatics, enterprise informatics, ecological informatics, medical informatics, security informatics, and social informatics. Such applications call for robust distributed infrastructure including databases, metadata, and knowledge bases, as well as tools for creation, deployment, and use of dis-

tributed data analysis pipelines that connect data sources, computational resources (for example, grid), and software services, that is, cyberinfrastructure for collaborative e-science. Similar challenges are shared by business, health-care, government, and national and global security applications.

The AAAI fall symposium on Semantic Web for Collaborative Knowledge Acquisition (SWeCKa-2006) was attended by approximately 30 researchers and practitioners with expertise and interests spanning artificial intelligence and semantic web, as well as several application areas. The 2.5 day symposium program consisted of invited and contributed talks as well as lively discussions among the symposium participants. Invited talks were given by Henry Lieberman of MIT Media Lab and Todd Hughes of DARPA. Lieberman spoke about collaborative construction of large commonsense knowledge bases by engaging large numbers of individuals to participate and contribute to such an exercise. Hughes spoke of research challenges and opportunities in security informatics. The contributed papers presented at the symposium included 11 regular papers, 4 short papers, and 6 extended abstracts, selected from a total of 40 submissions.

A majority of the presented papers covered one or more of the following topics: semantic integration of disparate information sources, collaborative construction and sharing of knowledge (including ontologies), collaborative knowledge acquisition from distributed data sources (such as on the web), and scientific, health-care, military, and social applications of such technologies. Some especially engaging discussions focused on research challenges in collaborative knowledge acquisition arising from the massive size, differences in query capabilities, semantics, and granularity of descriptions of data and knowledge sources. Others pointed out the need for shared benchmark data sets, ontologies, and mappings and open source software implementations and distributed test beds for empirical evaluation and further development of the different approaches. The symposium concluded with a plan of action for fo-

**The 2007 AAAI
Fall Symposium
Series will be held
November
8-11,
2007
at the Westin
Arlington Gateway,
Arlington Virginia,
adjacent to
Washington, DC**

cused workshops on related topics at major national and international conferences, and establishment of shared resources.

Vasant Honavar and Tim Finin cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-06 from AAAI Press.

Spacecraft Autonomy: Using AI to Expand Human Space Exploration

Autonomy will be a key component in future space exploration missions. Artificial intelligence technology can provide robust solutions to systems automation and mission operations coordination and can enable an incremental approach that extends from partial subsystem automation to coordination and oversight of distributed operations involving multiple systems and crew. The symposium had two goals: first to exchange information among a variety of efforts that all have the common thread of using AI to solve space-exploration problems; and second to define and understand the new challenges posed to the artificial intelligence community by new space-exploration projects.

The symposium was attended by



UW PROFESSIONAL MASTER'S IN COMPUTATIONAL LINGUISTICS

BLEND YOUR LANGUAGE AND PROGRAMMING INTERESTS INTO ONE WITH COMPUTATIONAL LINGUISTICS!

Earn your degree from one of the top linguistics programs in the nation. In only 12 months you can start a career in the exciting field of computational linguistics. Graduates go on to work as language software development engineers and linguistic data specialists for companies such as Adapx, AOL/Tegic, Google, Microsoft, the Palo Alto Research Center and Summer Institute of Linguistics.

www.compling.washington.edu



ICAIL 2007 4-8 June 2007
 11th Int'l Conference on
Artificial Intelligence and Law
 Stanford University, Palo Alto CA USA
<http://www.iaail.org/icail-2007>



ICAIL 2007 is the latest in a longstanding series of international conferences highlighting top quality research in AI and Law, with refereed papers and with proceedings published by ACM. The field is concerned with:

- the study of legal reasoning and argumentation using computational methods
- formal representation of norms and the role of norms in multi-agent systems
- investigation of AI technologies (e.g., machine learning, ontologies, tutoring systems) using law as an illustrative domain
- advanced legal IT applications (e.g., document generation, information retrieval, decision support, on-line dispute resolution)

This year's conference has special interest in Legal Applications for the *Semantic Web*.

ICAIL-2007 will be held under the auspices of the International Association for Artificial Intelligence and Law (IAAIL), an organization devoted to promoting research and development in AI and Law with members throughout the world. Additional sponsors include the Stanford Law School Program in Law, Science & Technology, and Thomson Legal & Regulatory.

Program Chair
Radboud Winkels, Leibniz Center for Law, University of Amsterdam, Netherlands
 (winkels@uva.nl)

Conference Chair
Anne Gardner, Atherton, CA USA (gardner@cs.stanford.edu)

more than 20 people drawn primarily from government and industry—both nationally and internationally. Matthew Barry, the Technical Integration Lead for Software and Avionics Interoperability and Reuse from NASA's Constellation Program Office, gave a keynote talk on "NASA's Exploration Software Requirements Influencing Autonomy and Automation," laying out high-level requirements for all of NASA's new exploration spacecraft. The goal is an evolvable software architecture that allows new technologies to be brought online as capabilities mature.

Another keynote address was given by Laura Forest and Lauren Kessler of Draper Laboratories. They are designing a new system to assist astronauts in landing their spacecraft on the moon. The talk, entitled "Human-Interactive Autonomous Flight Manager for Precision Lunar Landing," gave an overview of a software design that will reduce crew cognitive workload and

also reduce their dependence on earth-based ground controllers.

While on-board intelligence is an important component of spacecraft, there are also opportunities for artificial intelligence technology to assist ground controllers in their jobs. For every astronaut or robot in space, there are dozens of ground controllers involved in planning and monitoring the mission. Several talks were given on planning technologies that can automate routine activity planning that takes up a large part of ground control time. Robert Brown of Draper Laboratories gave a talk on a plan execution system (Timeliner) being used on the International Space Station (ISS) to automate some of the ground controllers' routine operations.

Robotics also received attention with several talks focusing on increasing their autonomy and independence. The next generation of rovers for both the Moon and Mars will have greater capabilities for on-board sci-

ence and on-board decision making. Robots stress the perception and action side of artificial intelligence.

Through all of the discussions, two common themes emerged. The first was the role of the human in any autonomy implementation. NASA human-rating requirements and the desires of crew and ground controllers require human involvement in system decision making. The term *adjustable autonomy* was often used to refer to the goal of maximizing the ability for human intervention even while minimizing its need. The second theme was validation and verification (V&V) of autonomous systems. The nature of many autonomy technologies will require new V&V techniques to be developed in order to compel trust in these technologies and their decisions.

Ari Jonsson and David Kortenkamp cochaired the symposium. Papers contributed to this symposium are available as Technical Report FS-06-07 from AAAI Press.