

# AAAI Fall Symposium Reports

*Jerry Ball, Chris Arney, Samuel G. Collins, Mitchell Marcus, Sergei Nirenburg, Antonio Chella, Kai Goebel, Jason Li, Margaret Lyell, Brian Magerko, Riccardo Manzotti, Clayton T. Morrison, Tim Oates, Mark Riedl, Goran P. Trajkovski, Walt Truszkowski, and Serdar Uckun*

■ The Association for the Advancement of Artificial Intelligence presented the 2007 Fall Symposium Series on Friday through Sunday, November 9–11, at the Westin Arlington Gateway, Arlington, Virginia. The titles of the seven symposia were (1) AI and Consciousness: Theoretical Foundations and Current Approaches, (2) Artificial Intelligence for Prognostics, (3) Cognitive Approaches to Natural Language Processing, (4) Computational Approaches to Representation Change during Learning and Development, (5) Emergent Agents and Socialities: Social and Organizational Aspects of Intelligence, (6) Intelligent Narrative Technologies, and (7) Regarding the “Intelligence” in Distributed Intelligent Systems.

## AI and Consciousness: Theoretical Foundations and Current Approaches

Is it possible to build a conscious machine? Is trying to design and build a conscious machine helpful to understanding the nature of consciousness? These questions have been at the core of AI since its beginnings. Recently there was an upsurge of interest in whether AI could address the issue of consciousness.

The primary goal of the symposium was to bring together experts from different disciplines who were interested in the relationships between AI and consciousness. The symposium was the first official place where scholars—coming from different fields as far as neuroscience and philosophy, psychology and computer science—addressed the issue of consciousness in a traditional AI environment. Furthermore, there was a good balance of American and European participants. The participants’ talks centered on the topic of the symposium and generated lively discussions of their research. Talks were balanced between implementation aspects and theoretical issues. The adopted relaxed scheduling let everybody present his or her own idea in a clear and pleasant way.

The symposium included two invited talks. The first talk, given by Giulio Tononi (University of Wisconsin–Madison School of Medicine, Depart-

ment of Psychiatry), focused on an insightful theory of integrated information as the foundation of conscious experience. The second talk, given by Aaron Sloman (University of Birmingham, School of Computer Science), summarized many crucial issues in the understanding of AI research inspired by the study of consciousness.

There was an almost generally accepted consensus among symposium participants on the following issues: (1) consciousness can be a source of inspiration for building better (more adaptive, more robust, more autonomous, more resilient) AI systems; (2) building AI systems aspiring to consciousness could be a testbed for psychological, philosophical, and neuroscientific theories of consciousness; (3) it is possible to address consciousness not only from neuroscience, psychology, and philosophy, but also from AI; and (4) the role of embodiment and situatedness is almost universally recognized.

A recurrent topic was the fact that the field of consciousness seems to be still far from a generalized consensus since it is rather unclear how to measure the advancements of various approaches, and further, it is often vague whether the work done really addresses the issue of consciousness or is just traditional AI.

One of the most crucial and controversial aspects was, not surprisingly, the dichotomy between phenomenal

and functional aspects of consciousness. Stevan Harnad (University of Southampton) painstakingly reminded the audience of the absence of any causal account of phenomenal aspects of consciousness and thus of the lack of any acceptable approach to feelings as such. In this respect, most symposium participants will accept the division between strong artificial consciousness (trying to build systems that are conscious) and weak artificial consciousness (trying to build AI systems that behave as if they were conscious or that clarify our understanding of consciousness).

In conclusion, consciousness (or the *c*-word as some ironically say) is a word that needs to be handled with care.

Antonio Chella (University of Palermo) and Riccardo Manzotti (IULM University) served as co-chairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-07-01.

— Antonio Chella and  
Riccardo Manzotti

## Artificial Intelligence for Prognostics

Over the last 10 years, there have been substantial interest and investment in prognostics in aerospace, transportation, and other industries. The field of prognostics focuses on methods and tools to determine functional degradation of components and systems and to estimate remaining useful life. The ultimate goal of prognostics is to manage the remaining useful life of systems such that maintenance actions can be performed “just in time” prior to failure, thus increasing safety as well as reducing maintenance expenses due to unscheduled downtime or unnecessary “preventive” maintenance. Given accurate remaining life estimates, prognostics also aims to manage the accumulation of further damage through control actions, for example, by either redistributing the load onto other components or changing the mission profile by trading off secondary mission goals.

These days, the hype around prognostics rivals the early days of artificial

intelligence. Nevertheless, in practice, accurate prognostics has proven rather difficult to accomplish. There are numerous issues that still need to be resolved before prognostics is adopted as standard practice in the industry.

The goal of the symposium was to explore how artificial intelligence and computational intelligence can aid in advancing this rapidly growing field. Contributions from researchers did focus on both data-driven and model-based approaches, with a bias toward data-driven approaches. Topics could roughly be partitioned into the areas prognostics methods, reconfiguration, decision-making, and uncertainty management.

Two major discussion topics emerged. One revolved around the lack of prognostic performance metrics and the other one on how to make prognostics actionable. Papers presented described different methodologies derived from math, physics, statistics, and of course artificial intelligence. But it became apparent that it was somewhat futile to derive any trusted conclusions based on the work presented. Analogies were drawn to a discipline where prognostics is integrated seamlessly into the workflow: medicine. Medicine has a wealth of prognostic tests and indicators that are used commonly for life-critical decisions. Examples include blood cholesterol, blood pressure, PSA, certain genetic markers for cancer, and so on. What makes it possible for medicine to trust these prognostic indicators is a common methodology: the randomized (double blind) clinical trial (RCT) combined with a statistical significance analysis (*z*-test). Arguably, the RCT is one of the critical advances in medicine that makes all progress possible. It was argued that what we’re missing in prognostics is the equivalent of an RCT. In the absence of that, there is a danger that everyone devises a performance methodology and reaches his or her own conclusions. In order to understand those conclusions, one needs to understand the performance methodology. That makes it nontrivial to compare the results. In order for prognostic health management (PHM) to become less of an art and more of a science, we will

need to develop a standard methodology similar to that of the RCT in medicine.

The second major discussion area revolved around the lack of known fielded prognostics systems, owing both to the novelty of the area and also to the need to overcome deployment hurdles. Business leaders need to make technology investment decisions based on perceived return on investment. Although conceptually prognostics is a game-changing technology, the quantification of the benefits is not straightforward. Indeed, prognostics is not performed for the sake of providing a remaining life estimate in itself but to serve as input for a decision engine that provides actionable decisions. Because it is the output from that decision engine that drives the benefit quantification, research in prognostics must also embrace the postprognostics decision making, in conjunction with uncertainty management, and validation and verification.

George Vachtsevanos (Georgia Institute of Technology), Serdar Uckun (NASA Ames Research Center), and Kai Goebel (NASA Ames Research Center) served as co-chairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-07-02.

— Kai Goebel and Serdar Uckun

## Cognitive Approaches to Natural Language Processing

This symposium highlighted research in natural language processing at the intersection of AI/computational linguistics and cognitive science/psycholinguistics. It also considered human-robot interaction (HRI) systems and communication with robots. The symposium followed the 2006 symposium “Between a Rock and a Hard Place, Cognitive Science Principles Meet AI Hard Programs” in exploring the interrelationship of AI and cognitive science within the context of natural language processing.

In his keynote address, John Trueswell (University of Pennsylvania) provided empirical evidence for con-

straint-based lexicalist theories of human language processing. That evidence supports parallel, interactive, constraint-based mechanisms operating over lexicalized chunks of linguistic knowledge and argues against serial, modular, and rule-based mechanisms operating over nonlexical representations. Jerry Ball (Air Force Research Laboratory) presented arguments for why it is important to consider well-established constraints on human language processing in the development of natural language-processing systems and why Trueswell's research is relevant and important for building functional natural language-processing systems.

In another keynote address, Alan Schultz (Navy Center for Applied Research in Artificial Intelligence) presented robotics systems that model humanlike behaviors including spoken language input and output, perspective taking, and determining frames of reference. Matthias Scheutz (Indiana University) presented a constraint-based, incremental natural language-processing architecture for developing situated embodied agents.

In a session on cognition and natural language processing, Gregory Aist (Arizona State University) suggested that there might be limits to how cognitively plausible we want natural language-processing systems to be. Do we want our systems to swear or make humanlike errors? Perhaps idealized cognition is more appropriate. Mitch Marcus (University of Pennsylvania) reiterated a theme for the symposium—schizophrenia! To what extent do we want to adhere to cognitive constraints on human language processing when trying to build functional natural language-processing systems? From Marcus's perspective the answer is "as cognitive as possible, but no more."

On day two, Marcus introduced a project to build a language-enabled robot capable of functioning as a team member in search and rescue. He emphasized that recent advances in the formal analysis of syntax, semantics and pragmatics, and statistical language processing make the goals of the project achievable. Among the advances in formal analysis is Aravind

Joshi's (University of Pennsylvania) lexicalized tree adjoining grammar (LTAG). LTAG makes use of lexicalized constituent trees, or supertags, which "complicate locally, but simplify globally." Christopher Potts (University of Massachusetts, Amherst) discussed the underspecification of meaning in language. Often meaning is not explicitly expressed but must be implicitly determined from the context and common ground.

A session on ontologies and agents began with the description of a project to build a simulated agent—the Maryland virtual patient—by Marjorie McShane and Sergei Nirenburg (University of Maryland Baltimore County). The virtual patient simulates the physiological progression of various diseases and the cognitive state of the patient. Norm Badler (University of Pennsylvania) presented his parameterized action representation (PAR) system for driving the behavior of virtual agents. Evgenia Malaia (Purdue University) presented an ontological framework using an event template for cross-linguistically representing the telic and atelic implications of "psych" verbs.

A session on the relevance of cognitive architectures and neuroscience for natural language processing began with a presentation by Wende Frost (Arizona State University) on modeling language within the Icarus cognitive architecture. Kevin Livingston (Northwestern University) argued that neuroscientific evidence provides support for the early application of semantic and episodic knowledge in language understanding in motivating an approach called direct memory access parsing. Andrea Heiberg (L3 Communication) showed how principles of optimality theory can be mapped into the ACT-R cognitive architecture in development of a language-generation capability. Nicholas Cassimatis (Office of Naval Research) argued for a language-processing approach that is grounded in a cognitive substrate that applies generally to all cognitive activity. Cassimatis was the first of two who provided insight into how to develop systems capable of integrating, constraining, and inferencing over representations across multiple levels of

representation. A second presentation on this subject was given by Andrew McCallum (University of Pennsylvania), who showed that solving the joint inference problem is key to moving away from staged models of language processing. Erwin Chan (University of Pennsylvania) presented a computational morphology-acquisition model that relies on unsupervised learning techniques. Unlike most statistical models, Chan's model performs better on sparse input.

The symposium concluded with a discussion of research directions and future steps. It is encouraging to see the recent emergence of funded research projects aimed at development of synthetic agents and robots capable of functioning as teammates in virtual and real-world environments. These projects refocus research on the deeper analysis of language required for full understanding of linguistic input. At the plenary session, Stuart Rodgers (Air Force Research Laboratory) summarized the results of the symposium. There has been and continues to be significant progress in natural language-processing research, and many of the researchers who attended the symposium are leading that research.

The symposium was organized by Chris Arney (Army Research Office), Mitch Marcus (University of Pennsylvania), Jerry Ball (Air Force Research Laboratory), Sergei Nirenburg (University of Maryland Baltimore County), and Marjorie McShane (University of Maryland Baltimore County). The papers from this symposium were not published.

— Jerry Ball, Chris Arney,  
Mitchell Marcu, and  
Sergei Nirenburg

## Computational Approaches to Representation Change during Learning and Development

Working with the right representation is important: a representation scheme determines what can be expressed and effectively computed. In machine-learning parlance, any representation

choice provides some inductive bias about what concepts can be learned with a given set of data. The role of representation is similarly understood in the fields of automated planning and scheduling, such as the relative merits of search in plan space versus state space. Besides managing representation change for artificial systems, cognitive and developmental psychology tells us that we (humans) routinely and facilely use different representations as we go about our day-to-day lives. Since at least Piaget's groundbreaking work in the 1930s, there has been the longstanding belief that the profound differences between newborn infants and adults is largely due to changes in the child's representation of the world. And beyond its importance at the level of individuals, the history of science is replete with breakthroughs that resulted from change of representations in scientific models, from the Copernican revolution in astronomy to Pasteur's germ theory of disease. These changes were not simply changes within scientific models, but large-scale reorganizations of our representation of the world, how observations are interpreted, and in some cases they changed the fundamental languages used to express models and the logic of science.

The primary goal of the symposium on Computational Approaches to Representation Change During Learning and Development was to understand how we can shift some of the burden of managing change of representation from ourselves to the systems we build. A secondary goal was to bring together researchers with different perspectives on the topic and to understand recent developments in their respective fields. Having computer scientists and cognitive and developmental psychologists in the same room for two and a half days generated invaluable discussion on the nature of representation change, what we know about representation change in humans, and how representation change processes might be turned into algorithms for automation. We discovered that there are a significant number of researchers studying representation change in its various forms

and that they have produced a number of concrete and successful approaches to the problem, though these researchers are spread across fields and do not yet form a unified scientific community.

It quickly became clear that there are different ideas about what counts as a change of representation. Everyone seemed to agree on the importance and relevance of discovering theoretical terms (or theoretical entities) causally efficacious though unobservable aspects of the environment (for example, black holes and, when they were posited by Pasteur, germs). Some prominent cognitive psychologists believe that adding theoretical terms (such as the object concept as discussed at length by Piaget) to one's representational repertoire is the only thing that truly counts as representation change. Others took less strident positions.

Josh Tenenbaum (MIT), for example, gave an invited talk on hierarchical Bayes methods for representation change, showing how to simultaneously represent the hypothesis space of models used to account for data, as well as a space of hypothesis languages in which those models are expressed.

A dimension of representation change that was a common theme of several presentations was that of feature learning, where features make up the fundamental units of a representation scheme. Approaches from machine learning combined feature learning with reinforcement learning methods to identify useful features that simultaneously improve learning. Another dimension of representation change reflected in both the psychological literature as well as AI knowledge systems is knowledge integration. Interesting work was reported in handling varying levels of granularity in knowledge representations, as well as approaches to learning and adapting ontologies.

Finally, work combining machine learning and robotics can now demonstrate learning representations for aspects of sensorimotor control that then become the representational primitives for a new level of learning, for example, about landmarks and

maps of the environment. Ben Kuipers (University of Texas) calls this kind of iterated representation learning "bootstrapping."

There was a clear sense at the end of the symposium that the time is right for the various researchers working on representation change in different fields to collaborate and develop a coherent subfield devoted to advancing the science of representation change.

Clayton T. Morrison (University of Southern California) and Tim Oates (University of Maryland Baltimore County) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-07-03.

— Clayton T. Morrison  
and Tim Oates

## Emergent Agents and Socialities: Social and Organizational Aspects of Intelligence

The study of agency and multiagent systems crosses disciplinary boundaries by focusing on society, culture, and communication 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. This symposium focused on second-order emergence. The constituents in a system are aware of an emergent phenomenon and adapt accordingly. New agents emerge as human and nonhuman agents interact, hinting at new qualities that may enable us to push the use of technology to its maximum capacity, and in the process imbricating both the observer and the observed in successive cycles of emergence. Theories of emergence suggest a dynamic, multidirectionality of perception organized socially as multiagent systems. What is less studied is the messiness of those multiagent systems themselves, the way they involve complex "translations" between human and

nonhuman agents or “transcodings” between different representational and discursive modalities.

The symposium brought together researchers from a variety of subfields of AI and other disciplines, such as philosophy and the social sciences. The symposium delved into the messiness of the social, approaching it from multiple perspectives simultaneously—computational, sociological, linguistic, and cybernetic—in such a way as to stimulate our own sites of emergence at the borders of these disciplines.

Three synchronized approaches on emergent phenomena dominated the symposium. One set of presentations focused on the definition of emergence, another on the multiagent simulation and modeling tools, and the third one on observations in artificial or biological societies.

The opening talk, given by Samuel Collins (Towson University), discussed the importance of the Macy Conference on emergence in hybrid multiagent systems and as such opened the discussion on the definitions of emergence. Marton Ivany (AITIA Inc.) presented the multiagent simulation suite and kicked off the discussion on the current and needed tools for simulation study of multiagent systems; Ivany was followed by the presentations of Yu Zhang (Trinity University) and Bogdan Werth (Centre for Policy Modeling, Manchester). David Newlin (TRI, Baltimore, MD) and his notes on the mirror neuron system started the discussion on emergence in biological societies.

Theodor Richardson (South University), Giulia Andrighetto and Marco Campenni (LABSS/ISTC-CNR, Rome, IT), and Giovanni Vincenti (Towson University) gave meaningful architectures for modeling artificial agents in artificial societies.

Goran P. Trajkovski (South University) and Samuel G. Collins (Towson University) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-07-04.

— Goran P. Trajkovski  
and Samuel G. Collins

## Intelligent Narrative Technologies

Research into artificial intelligence technologies that reason about, generate, and interactively manage narrative and story is motivated by the observation that narrative is a pervasive aspect of human culture; narrative plays a role in both entertainment and educational contexts of everyday life. As the prevalence of digital media technologies for both entertainment and education increases, the opportunity for innovative approaches to represent, perform, and adapt narrative experiences increases as well. The goal of the Intelligent Narrative Technologies symposium was to foster a community of researchers across disciplines that were interested in exploring computational approaches to narrative, story, and storytelling. The call for papers was intended to draw a wide net in terms of what was considered an intelligent narrative technology and also welcomed theoretical perspectives on narrative that could inform future intelligent narrative technologies. Fifty-six people attended the symposium, representing mostly computer scientists but also representing “new media” researchers and practitioners, narratologists, and fiction writers. While we would have liked a more diverse mix of disciplines, we were pleasantly surprised by the level of interest.

Twenty-six papers were presented (17 long-format papers and 9 short-format papers), which made for a busy schedule. A majority of the presentations were on interactive systems that used narrative to organize a user’s experience—often referred to as “interactive narrative systems.” We have witnessed a dramatic increase in the number of researchers investigating narrative in interactive, computer-mediated experiences as a device for more memorable and engaging context for users, likely due to the recent popularity of modern console and PC games and “serious games.” As an interesting counterpoint to the interactive narrative presentations, a half day’s worth of paper presentations was devoted to noninteractive perspectives on technologies such as story generation, sto-

ry understanding, and computational representations of narrative.

The symposium also included a discussion on authoring interactive narrative content, a highly interactive discussion on tangible storytelling interfaces, and a panel on the challenges of natural language generation for interactive narratives. The authoring discussion focused on authoring of narrative content, one of the bottlenecks to adoption of intelligent narrative technologies; unlike noninteractive narratives such as books and movies, interactivity results in a combinatorial explosion of a narrative content. An interesting series of relevant questions that arose were “How do we build these systems quickly?” and “What is the analogy between interactive narrative systems and the ‘indie’ short film?” Another interesting development at the symposium was the degree of interest in natural language generation. In retrospect, since narrative is by definition a form of discourse, it is more surprising that natural language generation has not been significantly represented in intelligent narrative technology research previously. The natural language generation panel overflowed its time boundaries into a lunchtime discussion.

The symposium concluded with a three-hour improvisational acting class lead by Brenda Harger of the Entertainment Technology Center at Carnegie Mellon University. The motivation for the class was to involve researchers in a real-world analogue to the kinds of experiences they are trying to represent computationally. Despite its being held after the official end of the last day of the symposium, 31 participants stayed for the class.

There had not been a symposium or conference in the United States devoted to narrative since the 1999 Fall Symposium on Narrative Intelligence. We feel that the symposium was an astounding success and a step toward solidifying an interdisciplinary community of researchers interested in intelligent computation and narrative. Talks are already under way on organizing follow-up events.

Brian Magerko (Michigan State University) and Mark Riedl (Institute for Creative Technologies, University of

Southern California) served as cochairs of this symposium. The papers of the symposium were published as AAAI Press Technical Report FS-07-05.

— Mark O. Riedl  
and Brian S. Magerko

## Regarding the “Intelligence” in Distributed Intelligent Systems

Looking across the range of application areas and application products today, the reoccurring keyword is “intelligence.” From web intelligence for business applications to coordinating robotic teams for NASA’s exploration vision to DOD’s net-centered approach to modern approaches for network security, all applications are expected to incorporate intelligence. The intelligence may be required for the application to succeed, or it may be an enhancement over a “dumber” version. Indeed, “intelligence” now serves as a system discriminator.

Themes of this symposium were to elucidate how the intelligence in a distributed system is expressed, to discuss which paradigms are proving particularly fruitful to support the expression of intelligence, and to look at what research areas are synergistic in supporting the expression of intelligence and which methodologies (and supporting tools) are useful for building intelligence into systems. Symposium participants included a mix of those from the academic, government, and private sectors.

Four invited speakers, over a two-day period, addressed several of the symposium themes. H. Van Parunak (TechTeam Government Solutions, Inc.) presented a talk titled “Monitoring and Managing Intelligence in Distributed Systems.” His talk discussed “swarm intelligence,” noting that intelligence may occur at the system level and is emergent intelligence, arising from interactions among entities (agents). The talk reviewed the application of concepts from statistical mechanics, those of entropy, chaos, and universality, to the tasks of monitoring and managing multiagent systems.

Katia Sycara’s (Carnegie Mellon University) talk, titled “Agents and Semantics in Open Distributed Environments,” introduced the semantic technologies and discussed their utility in developing open systems. The concept of “service” provided an additional organizing principle. Attention was given to the interaction of services and agents and the use of semantics to support interoperability. Sycara also presented architectural insight into developing such systems.

Jeffrey Bradshaw’s (IHMC) presentation on “Coordination in Human-Robot-Agent Teamwork” discussed the concept of “teamwork-centered autonomy” in facilitating interactions among humans, robots, and software agents. Requirements necessary for effective coordination among team participants and examples of policy design for coordination were presented. The talk culminated in a presentation of a human-agent-robot teamwork framework and its use in an operational exercise.

Lynne Parker’s (University of Tennessee, Knoxville) talk, titled “Distributed Intelligence: An Overview of the Field and its Application in Multirobot Systems,” provided an overarching view. Her talk categorized the types of interactions that may occur among entities in the distributed systems as being (1) collective, (2) cooperative, (3) collaborative, or (4) coordinative. Parker’s talk provided examples of how three now common “intelligence” paradigms, including the knowledge-based, ontological paradigm, the social or organizational paradigm, and the bioinspired paradigm, have been applied in systems involving multiple robots. She further discussed how the task-allocation problem for the multirobot systems is handled differently, depending on the paradigm that is employed.

A panel on Sunday focused on the impact of design on the “intelligence” in distributed intelligent systems. Panelists were Christopher Rouff (Lockheed Martin Advanced Technology Laboratories) and Wilbur Peng (Intelligent Automation, Inc.). The discussion was active and wide ranging, with features of and experiences with different distributed sys-

tems grounding the discussion. A consensus opinion is that the design strategies need to incorporate redundancy, control architecture, and a better awareness of overall operating conditions. The notion of “quality” in a distributed intelligent system was connected to the quality metrics for software systems, with an autonomic twist. The concept that an intelligent system should be able to tweak itself for better performance was also a consensus opinion.

A second theme to emerge from the panel discussion drew from the notion of a layered architecture applied to cognitive systems, with layers providing self-reflective, deliberative, and reactive capability. A major challenge in building distributed intelligent systems arises due to the difficulty in semantically integrating these layers.

Roughly 20 papers, given by participants over the course of the symposium, addressed aspects of intelligence in distributed intelligent systems. Many of the research papers addressed aspects of distributed intelligent systems that had been prototyped or deployed. The themes of individual papers roughly cluster into the areas that were defined by the invited talks.

The symposium concluded with a wrap-up session. Individuals provided their summary thoughts, many of which were echoed by multiple participants, and which included an interest in the wealth of domain areas and exemplars; the breadth and scope, from low-level to high-level intelligence; the diversity of the paradigms that may form hybrid systems and thought for how these paradigms may interact; and semantics, the use of which is a key to building distributed intelligent systems.

Walt Truszkowski (NASA Goddard Space Flight Center), Jason Li (Intelligent Automation Inc.), and Margaret Lyell (Intelligent Automation Inc.) served as chairs of the symposium. The papers of the Symposium were published as AAAI Press Technical Report FS-07-06.

— Margaret Lyell, Jason Li,  
and Walt Truszkowski