AAAI 2002 Fall Symposium Series Reports

Benjamin Bell, Lola D. Canamero, Silvia Coradeschi, Carla Gomes, Alessandro Saffiotti, Costas Tsatsoulis, and Toby Walsh

The American Association for Artificial Intelligence held its 2001 Fall Symposium Series November 2-4, 2001 at the Sea Crest Conference Center in North Falmouth, Massachusetts. The topics of the five symposia in the 2001 Fall Symposia Series were (1) Anchoring Symbols to Sensor Data in Single and Multiple Robot Systems, (2) Emotional and Intelligent II: The Tangled Knot of Social Cognition, (3) Intent Inference for Collaborative Tasks, (4) Negotiation Methods for Autonomous Cooperative Systems, and (5) Using Uncertainty within Computation. This article contains brief reports of those five symposia.

Anchoring Symbols to Sensor Data in Single- and Multiple-Robot Systems

The focus of this symposium was the connection between symbol- and signal-level representations of physical objects in autonomous robotic systems. We call *anchoring* the process of creating, and maintaining in time, this connection.

Anchoring must necessarily occur in any physically embedded system that comprises a symbolic reasoning component. A typical example is the problem of connecting, inside an autonomous robot, the symbol used by a planner to refer to a particular box, say *box-21,* to the data that correspond to this box in the sensorimotoric system. This connection must be dynamic because the same symbol must be associated to new entities in the perceptual stream to track the object over time or to reacquire it at a later moment. Anchoring must also occur in a multiple-robot system because the robots must agree about the meaning of the symbols used to refer to perceived objects in the environment.

Solutions to the anchoring problem are currently developed on a system-by-system basis for restricted domains. The aim of this symposium was to create an interdisciplinary community interested in the development of general theories of anchoring. Having such a theory will greatly advance our ability to build intelligent embedded systems and transfer techniques and results across different systems.

Attended by approximately 30 participants, the symposium was successful in achieving its aim. On the first day, each attendant was given two minutes to introduce his/her group and his/her research in a highly dynamic "rump session." The participants especially appreciated this session because many of them came from different communities and did not know each other. The remainder of the symposium was largely dedicated to the presentation of full-length papers; two invited talks; and even a virtual presentation, which covered different facets of the anchoring problem. Topics included the description of deployed robotic systems that use anchoring, the bottom-up extraction of symbols from sensor data, the anchoring of linguistic terms in human-robot or robot-robot communication, and related issues such as visual attention and conceptual spaces. Presentations differed significantly in the nature of the entities to be anchored: Some authors focused on symbols that denote individuals, and others focused on symbols that denote categories, actions, or events.

To further encourage creative interaction between the participants, the symposium included three panel discussions. Two short panels were devoted to the discussion of the relation between anchoring and other problems, including symbol grounding, pattern recognition, and tracking. A longer panel was organized at the end of the symposium and had the ambition to set up a research agenda for the anchoring problem. A crucial observation was that although the anchoring problem can easily become extraordinarily complex, we nonetheless have to deal with it if we want to build robots that manipulate their environment and interact with their users. There was general agreement that a first priority is to identify the right level of abstraction for a general theory of anchoring, which should be specific enough to lead to tractable solutions but still general enough that results and techniques can be ported across different robots and domains.

Besides the organized panels, a lot of spontaneous discussions arose during and after the presentations. Interestingly, a few recurrent themes emerged from these discussions, which appear to point to some fundamental issues for the anchoring problem. For example, it was noted that in most presented systems, anchoring relies on the existence of an internal representation for an object that is somehow intermediate between the symbol and the perceptual data. This intermediate representation plays a pivotal role in connecting different subsystems (for example, perception, reasoning, motion) that all need to refer to the same physical object. A second observation was that most presenters dealt with anchoring either bottom up (given some perception, associate it to a new or a preexisting symbol) or top down (given a symbol, find the corresponding data in the perceptual stream). It was felt that a general solution to anchoring should inte-

grate both directions. Finally, a question often popped up about whether we should focus on symbols that correspond to humanlike categories or symbols that correspond to categories closer to the robot's senses. Although the latter might be an easier starting point, we will eventually need the former to have our robots interact with humans in everyday life.

–*Silvia Coradeschi*
Alessandro Saffiotti
Orebro University

Emotional and Intelligent II: The Tangled Knot of Social Cognition

Three years after the first "tangled knot" meeting at the 1998 AAAI Fall Symposium in Orlando, Florida, this symposium gathered researchers from different disciplines and backgrounds to exchange ideas about the roles of emotions in grounding interpersonal behaviors and social cognition and to reflect on the motivations, scientific grounds, and practical consequences of our efforts to build artifacts endowed with emotional capabilities. Emotions seem to be at the heart of what being social means. This idea, however, leads somehow to a Gordian knot: Are we social because we have emotions, or do we have emotions because we are a highly social species? Should we rather talk of coevolution of emotions and sociality? Does this strong link between emotions, cognition, and sociality also apply to nonhuman, nonbiological species, that is, artifacts? These were some of the background questions underlying the different artifacts, models, and theories presented at the symposium.

Presentations were organized around five main themes: (1) emotion, social behavior, and learning; (2) perceiving of others' emotions and perceiving of others through emotions; (3) emotions in decision making; (4) creating and regulating of affect with and through artifacts; and (5) emotions in social behavior and adaptation.

Five keynote speakers addressed fundamental topics in various aspects of emotion research. The psychology of emotion was represented by Andrew Ortony, who presented "Some Thoughts about the Behavioral Concomitants of Emotions," and Craig A. Smith, who presented his work with Leslie Kirby entitled "On the Elicitation, Differentiation, and Organization of Emotion: Structural and Procedural Considerations." Rosalind Picard addressed a challenging problem in "Machine Recognition of Human Emotions"; John McCarthy expressed his concerns about emotional artifacts in "Robots Can Be Made with Humanlike Emotions but Shouldn't Be"; and Marvin Minsky presented some of the ground-breaking ideas about emotions he has been elaborating over the last years (and in his forthcoming book) in "The Emotion Machine: A Model of Everyday Mental Activity."

A panel discussion with the participation of Paul Dumouchel, John McCarthy, Marvin Minsky, Rosalind Picard, and Craig Smith was held the afternoon of the second day. It was entitled "The Impact of Affective Artifacts on Our Social World."

The symposium concluded with an open discussion where some big challenges and directions for future research in this area were identified. General consensus was reached about the benefits (or, rather, the need) of multidisciplinary efforts involving human, animal, and artificial emotion researchers. The role of AI and robotics in these efforts was also the object of agreement in that far from a mere implementation of psychological or biological theories about emotion, artificial models of emotions can provide unprecedented challenges and novel theoretical insights and have us think in different ways about issues relevant to emotions and their roles in social interaction and cognition. Looking toward the future, we see the spread of affective artifacts in our everyday life will also force us to rethink some of the ways in which we relate socially and emotionally with humans and machines.

–*Lola Cañamero*
Universite Paris-XI

Intent Inference for Collaborative Tasks

The symposium addressed issues surrounding *intent inference*, that is, observing a user's actions to generate inferences about the user's intent. Strong previous work in the area provided a foundation for the focus of this symposium, namely, scaling up intent inference to collaborative contexts, where team activity must be monitored and analyzed and inferences about team intention generated. Among our 32 participants were many of the pioneers of intent inference applications, 2 of whom offered keynote addresses to help provide a set of common reference points for the discussions. The group as a whole represented a breadth of disciplines, giving the conversations a decidedly multidisciplinary flavor that brought to the floor concepts from plan recognition, task modeling, work-flow analysis, probabilistic reasoning, natural language processing, and situated cognition, to name a few. Diversity was also pronounced in the affiliations of the participants, which reflected a balance of academic, industrial, and government organizations.

In her keynote talk, Christine Mitchell from the Georgia Institute of Technology provided a survey of a cluster of related research initiatives that use some form of activity tracking in generating inferences of user intention. Norm Geddes of Applied Systems Intelligence emphasized in his keynote talk the challenges of scaling up previous approaches to large-scale intent inference systems (for example, the NATIONAL AIR SPACE).

The presentations offered a noticeable breadth of ideas and examples; we learned about intent inference applications in process control, flight operations, command and control, and mobile commerce, for example. Investigators approached these problems from a range of disciplines and experimental methodologies, which enriched each participant's experience and kept the interchange lively.

Among the many areas of ongoing work our participants will continue to address are adversary intent inference, which shifts the modeling fo-

AIPS-2002—Tolouse, France

Proceedings of the Sixth International Conference on Artificial Intelligence Planning and Scheduling Systems

Edited by Malik Ghallab, Joachim Hertzberg, & Paolo Traverso

The 2002 proceedings papers span a wide range of topics, including papers that address real-world problems, bridge the gap between theory and practice, and effectively combine planning and scheduling with other areas of computer science. Many contributions focus on extending the state of the art and the rational foundations for critical issues such as interactive or mixed-initiative planning, dynamic planning, and uncertainty or incompleteness in planning.

ISBN 1-57735-142-8, $65.00 softcover

Available in July 2002

www.aaai.org/Press/

cus from the user or team to a hostile agent; intent-driven information retrieval and information extraction; and intent inference for mixed-initiative, human-agent collaboration.

By bringing together researchers from the intent inference community, those engaged in the study of collaboration, and prominent players in the application domains, this symposium succeeded in advancing an important dialogue. We have established a listserv for those interested in joining this discussion; visit www.coral.cs.cmu.edu/mailman/listinfo/agentmod-l to subscribe.

–Benjamin Bell
Lockheed Martin

Negotiation Methods for Autonomous Cooperative Systems

Negotiation is one important mechanism through which groups of autonomous systems can reach agreement, in a distributed fashion, on the sharing of limited resources or the allocation of tasks. Through negotiation, groups can form cooperative teams, in a bottom-up fashion, to resolve a variety of constraint-satisfaction problems. Such systems are typically characterized by (1) decentralized control, (2) partial and uncertain information, (3) and some sort of real-time constraint on resource or task allocation.

The symposium focused on three major topics over three days: The first day was dedicated to presentations and panels that discussed how negotiation could be used for the cooperative allocation of resources, especially in domains of soft or hard real-time constraints. The discussions centered on both theoretical and implementational issues of negotiation-based constraint satisfaction and resource and task allocation. Applications mentioned include sensor management for tracking and mission planning for unmanned aircraft.

The second day's presentations concentrated on theoretical issues and strategies of negotiation. Papers presented included studies of phase transition in games that model negotiation, ways of expressing and modeling preferences, and market models of negotiation. The panel discussion concentrated on how these theoretical approaches can be transitioned to more complex negotiation systems and protocols and how they can assist in the quality of the final solution achieved by the negotiating cooperative system.

The final half-day of the symposium looked at individualistic negotiation models and discussed whether and how these models could be applied to cooperative negotiation environment. The final panel (entitled "Negotiation versus the World: What's the Relationship of Negotiation Technology to Better-Known Technologies?") focused on why negotiation is needed and the types of problems where it can provide solutions that are better than those provided by other reasoning techniques. The panel raised more questions than it answered, and there was much discussion on the future of negotiation.

In addition to the paper presentations and the discussion panels, the symposium participants enjoyed two invited talks by Sarit Kraus and Sesh Murthy. Kraus concentrated on negotiation for dispute resolution, and Murthy presented various applica-

tions of negotiation to business-to-business environments.

–*Costas Tsatsoulis*
The University of Kansas

Using Uncertainty within Computation

We are witnessing an increasing interest in approaches based on randomization, probability, and uncertainty to speed up computation, model resources more realistically, and reason about complex systems. This symposium brought together researchers interested in this new area of research, bridging different areas of AI such as constraint programming, decision theory, flexible computation, knowledge representation and reasoning, learning, planning, probabilistic reasoning and probabilistic analysis of algorithms, qualitative reasoning, randomized algorithms, reasoning under uncertainty, and search.

There were four focus areas in the symposium: (1) algorithm selection and scheduling, (2) combinatorial search, (3) modeling uncertainty, and (4) applications. We were pleased to have Rina Dechter as an invited speaker. She gave a fascinating talk in which she explored the interface between constraint processing and Bayesian networks, between inference methods and probabilistic reasoning.

The methods used to boost search and deal with uncertainty discussed at the symposium ranged from the simple, for example, expectations and intervals, to the complex, for example, algorithm portfolios, Bayesian learning, and stochastic constraint programming. Search methods discussed included both local search and global (complete) search methods. The domains and applications presented at the symposium included randomized dispatch scheduling, reasoning across scenarios in planning under uncertainty, methods for sampling pages from the World Wide Web uniformly, temporal mechanisms for decision making in probabilistic networks, and stochastic context-free grammars.

The meeting included an interesting panel session entitled "Future of Uncertainty within Computation": Issues such as modeling and dealing with risk and trading off search by inference were featured extensively. It was recognized that this area of research is new, explicitly studying and developing methods based on randomization and uncertainty in general to boost our ability to model and reason about complex systems; nevertheless, we have only simple mechanisms that work, such as restarts within our algorithms. Our research agenda is more ambitious than that, and in the future, we will embrace uncertainty and build complex adaptive algorithms that cope with uncertainty. Overall, this is an exciting research area with tremendous potential for innovation in areas such as algorithm design and formal modeling of complex systems.

–*Carla Gomes*
Cornell University
–*Toby Walsh*
York University