AAAI 2007 Spring Symposium Series Reports

Thomas Barkowsky, Peter Bruza, Zachary Dodds, Oren Etzioni, George Ferguson, Piotr Gmytrasiewicz, Bernhard Hommel, Benjamin Kuipers, Rob Miller, Leora Morgenstern, Simon Parsons, Holger Schultheis, Adriana Tapus, and Neil Yorke-Smith

■ The 2007 Spring Symposium Series was held Monday through Wednesday, March 26–28, 2007, at Stanford University, California. The titles of the nine symposia in this symposium series were (1) Control Mechanisms for Spatial Knowledge Processing in Cognitive/Intelligent Systems, (2) Game Theoretic and Decision Theoretic Agents, (3) Intentions in Intelligent Systems, (4) Interaction Challenges for Artificial Assistants, (5) Logical Formalizations of Commonsense Reasoning, (6) Machine Reading, (7) Multidisciplinary Collaboration for Socially Assistive Robotics, (8) Quantum Interaction, and (9) Robots and Robot Venues: Resources for AI Education.

Control Mechanisms for Spatial Knowledge Processing in Cognitive / Intelligent Systems

The purpose of this symposium was to address and investigate the interface and possible interplay between spatial knowledge processing and control processes. The former refers to the coding and use of spatial information in the perception of, the navigation in, and the communication about spatial configurations. The latter refers to all those processes that organize and integrate information, allocate processing resources, and tailor information streams to the current conditions so as to allow for coherent functioning of biological and artificial cognitive systems in their environment.

Although both areas have been researched intensely in the past, the question of how they interface has received only little attention. One reason for this could be that the two areas are modular and orthogonal to each other, so they can be readily investigated separately. And yet, essential characteristics of their interaction may be identifiable only when they are investigated together.

Either way, only a little research has explicitly addressed both this central question and the precise form of the control mechanisms involved in spatial knowledge processing. The main aim of the symposium was to bring together researchers to broach explicitly the issue of control in spatial knowledge processing for the first time.

We had a highly interactive symposium with contributions from people representing a wide range of disciplines: artificial intelligence, cognitive psychology, linguistics, neuroscience, and cognitive robotics. The symposium was thematically structured by a number of selected presentations. The format of the symposium combined short plenary presentation sessions with small topical breakout sessions (in parallel) followed by plenary report-back cycles. In this way, all participants were actively involved a considerable amount of the time. The main emphasis was on producing and exchanging new ideas, perspectives, and topics for further research.

In the scope of the symposium, several ways of implementing control mechanisms for spatial knowledge processing were proposed, from an AI (for example, case-based reasoning), a robotics (such as reproductive perception), and a cognitive-modeling (for example, modeling the central executive) perspective. Regarding the more fundamental issue mentioned above, some participants of the symposium were advocating the view that spatial knowledge processing is controlled as any other kind of information processing, and thus, cognitive architectures such as Soar or ACT-R constitute sufficient frameworks for modeling spatial knowledge processing. Other participants, however, deemed control in spatial knowledge processing (for example, the coordination, combination, and integration of multiple spatial representations) special and not covered by existing, general-control mechanisms.

Besides these more specific results and approaches, the symposium as a whole revealed that research concerning control in spatial knowledge processing is still in its infancy. For instance, central concepts are only vaguely defined and are in need of clarification. The momentum induced by this event is expected to promote research activities toward gaining a deeper understanding of how control mechanisms for spatial knowledge processing are or should be realized in natural and artificial cognitive systems, respectively.

The papers from this symposium were published in the AAAI technical

  ISSN 0738-4602

report series and are available from AAAI Press.

—Holger Schultheis, Thomas Barkowsky, Benjamin Kuipers, and Bernhard Hommel

Game-Theoretic and Decision-Theoretic Agents

This symposium marked the tenth in a series of successful game theory and decision theory symposia and workshops held over the last 10 years. The symposium attracted submissions and participation of researchers interested in principled techniques of decision and game theories to design autonomous agents. Decision theory provides a general paradigm for designing rational agents capable of operating in partially observable and nondeterministic environments. Decision-theoretic models use precise mathematical formalism to define the properties of the agent's environment, the agent's sensory capabilities, the ways the agent's actions change the state of the environment, and the agent's goals and preferences. The agent's rationality is defined as behavior that maximizes the expectation of the degree to which the preferences are achieved over time, and the planning problem is identified as a search for the optimal plan.

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

This year's submissions reflected the wide range of topics in planning, interacting, and learning. The discussions centered on the complementary ways the techniques of decision theory and game theory should be used to obtain designs of competent agents, learning to achieve efficient interaction or equilibria, approximate solutions to optimal planning problems, novel solution concepts, and automatic mechanism design. One of the highlights was an invited talk by Hal Varian, from the Hass Business School at the University of California at Berkeley, which illustrated the benefits of equilibrium analysis in designing auctions for position of advertisements appearing on Google.

The papers from this symposium were published in the AAAI technical report series and are available from AAAI Press.

—Piotr Gmytrasiewicz and Simon Parsons

Intentions in Intelligent Systems

Intentions, in the sense of agents having specific purposes in mind when they do things, have long played a central and organizing role in the analysis of intelligent behavior. The AAAI 2007 Spring Symposium on Intentions in Intelligent Systems considered the role of intentions in implemented (or reasonably foreseeable) AI systems. The focus was primarily on practical, realistic systems that perform tasks intelligently, as opposed to abstract philosophical theories of intention or purely mathematical formalisms for representing intention. The symposium brought together key