

## INVITED TALKS

### TOWARD AUTOMATING SCIENTIFIC DISCOVERY

Bruce G. Buchanan, *Professor of Computer Science, University of Pittsburgh*

Discovery in science is often described in such terms as serendipity, insight, creative impulse, and other attributes thought to be unmechanizable. The process of science, however, is much like other problem-solving processes in which large problems can be broken into subproblems and experienced scientists with more knowledge can be seen to be better than novices. In this talk I revisit several of the critical steps in mechanizing scientific discovery and describe prototype AI programs from our laboratory that address those steps.

### MEANING FROM DATA: CHALLENGES FOR ROBOT BABY

Paul Cohen, *Professor of Computer Science, University of Massachusetts, Amherst*

When you program a computer, you make it process symbols in such a way that the results mean something to you. Whether the symbols and results mean anything to the computer is not a consideration to programmers, nor is it an easy philosophical question. Yet it is important because computers cannot catch semantic errors unless they know the meanings of symbols. In the Robot Baby project we have robots learn the meanings of symbols. I will describe work on learning the meanings of words in speech and text, and the meanings of states in planning. As a class, the learning algorithms find structure in multivariate time series, and they are generally useful in data mining from sequences and series. We have been interested to see that different aspects of meaning are extracted by different algorithms. I will conclude by describing two challenge problems that have valuable applications and also drive research in the direction of semantic autonomy for agents.

### COLLABORATIVE SYSTEMS: PROSPECTS AND PROBLEMS

Barbara J. Grosz, *Higgins Professor of Natural Sciences  
Division of Engineering and Applied Sciences, Harvard University*

As a result of the ubiquity of computer networks, computer systems increasingly participate in complex, distributed communities of people and systems, rather than operating as solitary devices employed by a single person. This major shift in the way we use computers has led to a significant challenge for computer science, one in which AI can play a major role: to determine ways to construct computer systems that are able to act effectively as collaborative team members. Teams may consist solely of computer agents, but often include both systems and people. They may persist over long periods of time, form spontaneously for a single group activity, or come together repeatedly. This talk will briefly review the major features of one model of collaborative planning, SharedPlans (Grosz and Kraus, 1996, 1999) and will describe efforts to develop collaborative planning agents and systems for human-computer communication based on this model. The model also provides a framework in which to raise and address fundamental questions about collaboration and the construction of collaboration-capable agents. The talk will describe recent approaches to three such problems—commitment management, assessment of alternative courses of action, and group decision-making for recipe selection.