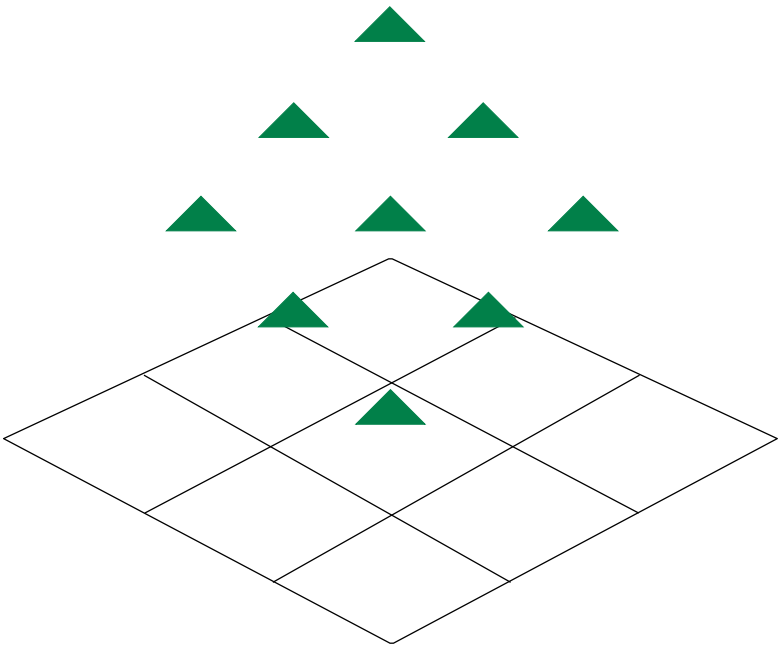

AAAI 1996

Fall Symposium Series

November 9 – 11, 1996
Massachusetts Institute of Technology
Cambridge, Massachusetts

Call for Participation

Sponsored by the
American Association for Artificial Intelligence
445 Burgess Drive, Menlo Park, CA 94025
(415) 328-3123
fss@aaai.org
http://www.aaai.org/Symposia/symposia.html

The American Association for Artificial Intelligence presents the 1996 Fall Symposium Series, to be held Saturday through Monday, November 9-11, 1996 at the Massachusetts Institute of Technology. The topics of the seven symposia in the 1996 Fall Symposium Series are:

- Configuration
- Developing Assistive Technology for People with Disabilities
- Embodied Cognition and Action
- Flexible Computation in Intelligent Systems: Results, Issues, and Opportunities
- Knowledge Representation Systems Based on Natural Language
- Learning Complex Behaviors in Adaptive Intelligent Systems
- Plan Execution: Problems and Issues

Symposia will be limited to between forty and sixty participants. Each participant will be expected to attend a single symposium. Working notes will be prepared and distributed to participants in each symposium.

A general plenary session, in which the highlights of each symposium will be presented, will be held on Sunday, November 10. An informal reception will be held on Saturday, November 9.

In addition to invited participants, a limited number of other interested parties will be able to register in each symposium on a first-come, first-served basis. Registration information will be available in early August. To obtain registration information, write to:

AAAI
Fall Symposium Series
445 Burgess Drive
Menlo Park, CA 94025

Information can also be obtained by email (fss@aaai.org) or on AAAI's web page at:

http://www.aaai.org/Symposia/symposia.html.

Submission Dates

- Submissions for the symposia are due on April 15, 1996
- Notification of acceptance will be given by May 17, 1996
- Material to be included in the working notes of the symposium must be received by August 23, 1996.

See the appropriate section below for specific submission requirements for each symposium.

This document is available as http://www.aaai.org/Symposia/Fall/1996/fall-participation.html

Configuration

Configuration involves selecting and arranging parts to fit problem constraints. Configuration problems may involve design, manufacturing, sales, installation and maintenance. The parts need not be physical, e.g. they may be actions in plans. Configuration problems have a long history in AI going back at least to the pioneering R1/XCON expert system for configuring computer systems. Recently the area has been revitalized by:

- Renewed industrial interest: Edward Feigenbaum highlighted configuration in his "Tiger in a Cage" talk at AAAI-93.
- The rise of constraint satisfaction technology: Constraint programming languages are bringing this technology into the real world. Configuration is a natural constraint satisfaction problem.
- The success of the world wide web and other networked information services: These generate an increasing demand for automated configuration. Companies are seeing their offers/sales ratio increase dramatically to the point where manual configuration is no longer feasible.

At present, academic attention in AI to this topic is lagging somewhat behind the industrial level of interest. The symposium aims to bring together academia and industry to advance a scientific framework for configuration research informed by the demands of practical configuration problems.

Current information on the symposium will be maintained at: http://liawww.epfl.ch/aaai-fs96/CFP.html

Submissions

Potential participants should submit:

- Name, physical and electronic addresses, also fax number and WWW URL if available. If several people working together, e.g. collaborating authors, wish to attend, each should submit separately, but should also name the others in the group.
- A brief statement describing why you wish to attend and how you believe that you can contribute to the symposium. List related work that you have done, papers you have written or programs you have developed. Can you contribute a demo, a video, a benchmark problem, a brainstorming exercise?
- A paper for the working notes, no longer than 8 pages. This may be a research paper, a description of a working system or practical problem domain, or a position paper.

Send both hard copy and electronic versions of this material to:

Rainer Weigel
AI Laboratory (EPFL)
Swiss Federal Institute of Technology
IN-Ecublens, CH-1015
Lausanne, Switzerland
configuration-ws@lia.di.epfl.ch

Organizing Committee

Boi Faltings (Cochair), Swiss Federal Institute of Technology, faltings@di.epfl.ch; Eugene C. Freuder (Cochair), University of New Hampshire, ecf@cs.unh.edu; Alois Haselboeck, Siemens AG, hasel@garwein.hai.siemens.co.at; Ken MacCallum, University of Strathclyde, ken@cad.strath.ac.uk; Deborah L. McGuinness, AT&T Bell Laboratories, dlm@research.att.com; Sanjay Mittal, Catalogics Software Corporation, mittal@ix.netcom.com.

Developing Assistive Technology for People with Disabilities

There is a growing interest in applying the theory and techniques developed by AI research to the domain of assistive technology for people with disabilities and in developing new research within the assistive technology domain. Some areas of current work include the development of new user interfaces for computers to accommodate people with varying motor, hearing or sight disabilities, robotic wheelchairs, speech recognition systems for people with hearing disabilities, text to speech systems for blind people, and automation of the process of converting textbooks and other written materials into recordings for the blind.

This application domain is particularly interesting because the interaction between the person and the system allows researchers to overcome some of the common stumbling blocks for AI applications. Normally, AI applications attempt to solve all possible situations in a domain. Assistive applications are intended to work in conjunction with a person with limited vision, hearing or motor capabilities. Therefore, assistive applications need only solve a portion of the problem, while leaving unsolved aspects to the user.

While the addition of a person into the cognitive loop allows researchers to avoid some of the usual difficulties, it adds a new dimension that must be considered: the user interface. Researchers must consider the needs of people with disabilities, often including interviews in the research process. Assistive applications with ineffective user interfaces are useless. Research in this area needs to integrate AI technology with UI technology to come up with new solutions. Our goals are to:

- Initiate a dialogue between the AI community and other research communities that will facilitate an exchange of ideas to further research progress.
- Identify areas of AI research that can be used to solve problems in this domain.
- Discuss how research in this domain can further general research in AI.
- Present successfully implemented systems.
- Discuss how to address user interface issues when designing systems.

Submissions

Potential participants should submit a short paper (5-8 pages) describing work in progress or completed work. Other interested participants should submit a one to two page description of their work in this area (including a short list of related publications) or specific questions and issues that they feel should be addressed. Please send your submission(ASCII files preferred) via e-mail to holly@ai.mit.edu. Accepted participants will be asked to submit PostScript versions of their papers. People interested in demonstrating their system or showing videotapes (either in addition to or in lieu of a paper presentation) should contact Holly Yanco at holly@ai.mit.edu prior to the deadline.

Organizing Committee

Holly Yanco (Chair), MIT, holly@ai.mit.edu; John Aronis, Univ. of Pittsburgh, aronis@cs.pitt.edu; David Miller, KISS Institute for Practical Robotics, kipr@src.umd.edu; Vibhu Mittal, Univ. of Pittsburgh, mittal@cs.pitt.edu; T.V. Raman, Adobe Systems, raman@adobe.com.

Embodied Cognition and Action

The role of physical embodiment in cognition has long been the subject of debate. It is largely accepted in AI that embodiment has strong implications on the control strategies for generating purposive and intelligent behavior in the world. Some theories have proposed that embodiment not only constrains but may also facilitate certain types of higher-level cognition. Evidence from neuroscience allows for postulating shared mechanisms for low-level control of embodied action (e.g., motor plans for limb movement) and higher-level cognition (e.g., abstract plans). Work in animal behavior has also addressed the potential links between the two systems and linguistic theories have long recognized the role of physical and spatial metaphors in language.

The symposium will study the role of embodiment in both scaling up control and grounding cognition. We will explore ways of extending the existing typically low-level sub-cognitive systems such as autonomous robots and agents, as well as grounding more abstract typically disembodied cognitive systems. We will draw from AI, ethology, neuroscience, and other sources in order to focus on the implications of embodiment in cognition and action, and explore work that has been done in the areas of applying physical metaphors to more abstract higher-level cognition.

Topics and questions of interest include:

- What spatial metaphors that can be used for abstract/higher-level cognition?
- What non-spatial metaphors can be applied in higher-level cognition?
- What alternatives to symbolic representations (e.g., analogical, procedural, etc.) can be successfully employed in embodied cognition?
- How can evidence from neuroscience and ethology benefit work in synthetic embodied cognition and embodied AI? Can we gain more than just inspiration from biological data in this area? Are there specific constraints and/or mechanisms we can usefully model?
- (How) Do methods for modeling embodied insect and animal behavior scale up to higher-level cognition?
- How do metaphors from embodiment apply to everyday activity?
- What computational and representational structures are necessary and/or sufficient for enabling embodied cognition?
- What are some successfully implemented embodied cognition systems?

The symposium will focus on group discussions and panels with a few inspiring presentations and overviews of relevant work.

Submissions

We invite the participation of researchers who have been working on embodied cognition and action in the fields of AI, neuroscience, ethology, and robotics.

Prospective participants should submit a brief paper (5 pages or less) or an extended abstract describing their research or interests. Papers should be submitted electronically via ftp to ftp.cs.brandeis.edu/pub/faculty/maja/aaai96-fs/. Participants will have an opportunity to contribute to the final working notes.

The WWW home page for this symposium can be found at: http://www.cs.brandeis.edu /~maja/aaai96-fs

Organizing Committee

Maja Mataric (chair), Brandeis University, maja@cs.brandeis.edu; Dana Ballard, University of Rochester, dana@cs.rochester.edu; Rod Brooks, MIT, brooks@ai.mit.edu; Daniel Dennett, Tufts University, ddennett@pearl.tufts.edu; Simon Giszter, Medical College of Pennsylvania, simon@SwampThing.medcolpa.edu; Erich Prem, Austrian AI Institute, erich@ai.univie.ac.at; Terence Sanger, MIT, tds@ai.mit.edu; Stefan Schaal, Georgia Tech, sschaal@cc.gatech.edu

Flexible Computation in Intelligent Systems: Results, Issues, and Opportunities

Flexible computation refers to procedures that allow a graceful tradeoff between the quality of results and allocations of costly resources, such as time, memory, or information. Systems employing flexible computation gain the ability to adapt the quality of their response to dynamic changes in requirements for precision, and to uncertainty or variation in the cost of computational commodities. Recent examples of flexible computation techniques include memory-bounded search, anytime algorithms, approximate query processing, and a variety of imprecise computation techniques.

Flexible computation has been applied to combinatorial optimization, planning, probabilistic inference, decision making, and theorem proving. Our goal is to explore results, critical problems, and opportunities via invited talks, presentation papers, and panel discussions.

Topics of interest include:

- Fundamental properties of flexible computation
- Scheduling and monitoring of flexible computation
- Partitioning resources between object-level and meta-level
- Representation and learning of performance profiles
- Flexible system specification and evaluation
- Programming techniques supporting construction and composition
- Benefits and overhead associated with use of flexible computation
- Applications of flexible computation

We are especially interested in experiences with applications in time-critical environments, or other resource-limited situations.

Submissions

We invite the participation of researchers who have been working on problems in areas including artificial intelligence, information retrieval, databases, operating systems, communications, signal processing, robotics, and numerical analysis. We are interested in both analytical and experimental work.

Prospective participants should submit a brief paper (5 pages or less) or an extended abstract describing their research or interests. Papers should be submitted electronically via ftp to flexcomp.microsoft.com /flexcomp96. Participants will have an opportunity to contribute to the final working notes.

Further information can be found on the WWW home page for this symposium: http://flexcomp.microsoft.com

Organizing Committee

Eric Horvitz (Cochair), Decision Theory Group, Microsoft Research, horvitz@microsoft.com; Shlomo Zilberstein (Cochair), University of Massachusetts, shlomo@cs.umass.edu; Louis Hoebel, Rome Laboratory/C3CA, hoebel@ai.rl.af.mil; Jane Liu, University of Illinois, janeliu@cs.uiuc.edu; Mike Pittarelli, SUNY Institute of Technology, mike@sunyit.edu .

Knowledge Representation Systems Based on Natural Language

Note: This Symposium will follow KR-96 (to be held November 5-8, 1996)

The Symposium addresses the theoretically and practically important problem of knowledge representation (KR) systems that closely parallel the representational and inferential characteristics of natural language (NL).

Advantages of such NL-based KR systems would be enormous. Among the arguments for the natural-language-as-KR-system approach are:

- KR systems based on natural language would be easy for people to use,
- Most human knowledge is encoded and communicated via natural language, in the form of textual documents and (transcribed) interactions (dialogs). A NL-based KR system would be capable of automatically creating and updating its knowledge base from natural language texts more easily. Additionally, the contents of this knowledge base and inferences supported by the KR system would parallel those of a natural language user.
- Every day, a huge number of new textual documents becomes available on-line. This creates the need for more sophisticated information retrieval techniques based on natural language processing (NLP) and KR techniques.
- KR systems based on natural language would provide a uniform symbolic representation. The same representational and inference mechanism could be used when utilizing previous knowledge for processing new natural language inputs (natural language as both meta-level and object-level language).
- It is hard to match expressiveness and precision of natural language, particularly in not (well) formalized domains.
- Many philosophers, linguists and cognitive scientists believe that mental-level representation of knowledge (human mind) is close in form to natural language.

While some AI researchers believe that it is feasible and necessary to design KR systems closely mimicking natural language, others are pessimistic about success or even possibility of designing such KR systems. This pessimism might account for the general lack of interest in the problems of NLP within the KR community; for example, only six of the twenty-two KR systems presented in the "Special Issue on Implemented Knowledge Representation and Reasoning Systems," Charles Rich, Editor *SIGART Bulletin,* Vol. 2 (3), ACM Press, 1991, are driven by NLP concerns.

Among the arguments against the NL-as-KR-system approach are:

- Natural language is (highly) ambiguous,
- Natural language has (very) complex syntax, semantics, and pragmatics,
- Natural language is non-systematic, non-algorithmic,
- Natural language is (highly) context-dependent,
- Natural language is (merely) an interface. Inferencing does not belong with natural language.

The goal of this symposium is to address in-depth such arguments for and against designing KR systems closely simulating natural language.

We invite papers that substantiate the view that natural language can be viewed as a KR system with

its own representational and inferential machinery, and, as such, is a productive source of ideas for KR formalisms and their practical implementations.

We are interested in papers discussing representations and inference mechanisms paralleling a nontrivial or interesting subset of natural language and formal systems whose expressiveness, semantics, information packaging, reasoning, and computational tractability closely correspond to that of natural language.

We are interested in automatic or semi-automatic methods of obtaining taxonomies facilitating various NLP tasks such as anaphora resolution, inferencing, and machine translation.

We are also interested in papers that discuss those aspects of natural language that are not desirable in a KR system. We invite position papers with supported arguments against the idea of designing KR systems that mimic natural language.

Submissions

Paper submissions should be at *most* 15 pages (including title, abstract, figures, but excluding references) in 12 pt font article LaTex style. The first page must include title, author's name(s), affiliation, complete mailing address, e-mail address, phone/fax number(s), abstract of 200 or so words, and keywords.

As a *last* resort, five hard copies of the paper can be mailed to:
Lucja Iwanska
Department of Computer Science
Wayne State University
Detroit, MI 48202, USA
(313) 577-1667 (phone)
(313) 577-2478 (secretary)
(313) 577-6868 (fax)

Papers should be submitted electronically via ftp to ftp.cs.wayne.edu/pub/nlkr. You are encouraged to e-mail your intention to submit by January 1, 1996 to lucja@cs.wayne.edu.

Further information on this symposium will be available at URL http://www.cs.wayne.edu/nlkr

Organizing Committee

Lucja Iwanska (Chair), Wayne State University, lucja@cs.wayne.edu; Syed S. Ali, Southwest Missouri State University, syali@sy.smsu.edu; Douglas Appelt, SRI International; appelt@ai.sri.com; R.V. Guha, Apple Computers, Inc., guha@taurus.apple.com; Sasa Buvac, Stanford University, buvac@sail.stanford.edu; Douglas Lenat, CYC Corp., lenat@mcc.com; David McAllester, AT&T Bell Labs, dmac@research.att.com; Len Schubert, University of Rochester, schubert@cs.rochester.edu; Stuart C. Shapiro, State University of New York at Buffalo, shapiro@cs.buffalo.edu; Wlodek Zadrozny, IBM TJ Watson Research Center, wlodz@watson.ibm.com.

Learning Complex Behaviors in Adaptive Intelligent Systems

The machine learning community made an important methodological transition by identifying a collection of benchmarks that can be used for comparative testing of learning algorithms. While the resulting comparative research contributed substantially to progress in the field, a number of recent studies have shown that very simple representations such as depth-two decision trees or perceptrons perform relatively well on many of the benchmarks (which are typically static fixed-size databases). At the same time, the hand-crafted knowledge representations for solving complex tasks are typically rather large and are often designed to cope with complex dynamic environments.

This symposium will attempt to bridge this gap by focusing on algorithms that learn to perform complex behaviors and cognitive tasks, such as reasoning and planning with uncertainty, perception, natural language processing and large-scale industrial applications. The scalability of learning systems, such as reinforcement learning, is therefore a principal focus of the symposium. The underlying theme is the automated construction and improvement of complete intelligent agents, which is closer in spirit to the goals of AI than learning simple classifiers. We expect to have an interdisciplinary meeting with participation of researchers from AI, neural networks, machine learning, uncertainty in AI and computer science theory.

Some of the key issues we plan to address are:

- Research on agents that learn to behave "rationally" in complex environments.
- Generating new benchmarks and devising a methodological framework for studying empirical scalability of algorithms that learn complex behaviors.
- Measures of complexity of learning problems involving ongoing behavior in a dynamic environment.
- Empirical and theoretical analysis of the scalability of different representations and learning methods.
- Development of new theoretical frameworks for analysis of broader learning tasks such as learning to reason, learning to act, and reinforcement learning.

The symposium will consist of invited talks, submitted papers, and panel discussions on the above issues and any related topics. We will invite short technical papers on these issues as well as position papers relating learning and issues in knowledge representation; comparative papers that illustrate the capabilities of different representations to achieve the same functionality; and papers providing specific benchmarks that demonstrate the scalability of a particular representation or paradigm.

Submissions

Prospective participants are encouraged to submit extended abstracts (5-8 pages) addressing the research issues above. Please refer to an extended version of the call for papers that provides additional submission information and a tentative program (available on the WEB at: http://www.cs.jhu.edu/complex/symposium/cfp.html). Electronic submissions as well as inquiries about the program should be sent to complex@cs.jhu.edu.

Organizing Committee

S. Kasif (Cochair), Johns Hopkins University, kasif@cs.jhu.edu; S. Russell (Cochair), Berkeley, russell@cs.berkeley.edu; B. Berwick, MIT, berwick@ai.mit.edu; T. Dean, Brown University, tld@cs.brown.edu; R. Greiner, Siemens Research, greiner@scr.siemens.com; M. Jordan, MIT, jordan@psyche.mit.edu; L. Kaebling, Brown University, lpk@cs.brown.edu; D. Koller, Stanford University, daphne@cs.berkeley.edu; A. Moore, CMU, awm@cs.cmu.edu; D. Roth, Weizmann Institute, danr@das.harvard.edu.

Plan Execution: Problems and Issues

Traditionally, work on planning has concentrated on how plans can be constructed, but as the planning community has begun to address more complex, real-world problems, the issues surrounding the execution of plans have come to the fore. Systems are being built that construct plans that are used in domains ranging from robot navigation to image processing and information retrieval. Such systems must often be able to handle actions with duration, simultaneous execution of actions, plans with conditionals and loops, and plan failure. These new demands on an agent require considerable extensions to the classical model of plans as simple sequences of actions guaranteed to achieve their goals.

In this symposium we shall address problems associated with executing plans in real-world domains. Such domains have many, if not all, of the following characteristics:

- Complexity: neither the plan construction system nor the plan execution system (if different) can have complete information;
- Dynamism: the world can change independently of the plan being executed (whether through the actions of other agents or through exogenous events);
- Uncertainty: the results of performing an action often cannot be predicted with certainty;
- Interruptibility: actions may last over appreciable durations, and may be interrupted during their execution;
- Concurrency: actions and events may occur simultaneously;
- Changing Objectives: new goals can arise and old goals can become unimportant as time passes;
- Goal Variability: goals may vary along a spectrum from maintenance goals (keep the value of G as close to V as you can) to achievement goals (make G true then terminate).

We shall look at the plan execution problem from the point of view of the system executing the plan by considering two interdependent aspects: the nature of executable plans and how they should be executed.

First, does the nature of a plan change when one takes into account the possibility of feedback, failure, and recovery? How should such a plan be represented to facilitate its execution? What information should it include? Should plans be sketchy or detailed? An important issue is how the answers to these questions vary with the capabilities of the system executing the plan.

Second, we shall consider how plans should be executed. How should plan failure be recognized and how should the recovery or repair process proceed? Unexpected changes in the world may be either adverse or benevolent: how can a system exploit run-time opportunities to improve its performance? Once again, how does the approach to these issues depend on the form that the plans take?

Submissions

Potential participants should submit either an extended abstract of up to 5000 words or a position paper of up to 2 pages.

All extended abstracts, whether describing a working system or more theoretical in nature, should answer the following questions:

- Which of the domain characteristics listed above were considered especially significant in the work described?
- Which other characteristics does your domain have, or does your theory account for?
- What are the limitations of your system or theory? Which of the domain characteristics would cause problems for it?

In addition, extended abstracts should include worked examples set in a simulated or actual real-world domain.

Position papers should summarize an approach to an issue or issues in plan execution, explaining which domain characteristics are considered especially relevant.

Electronic submissions only will be accepted. Email ASCII (position papers only) or PostScript files to the symposium chair, Louise Pryor louisep@aisb.ed.ac.uk.

Further information on this symposium will be available at URL http://www.dai.ed.ac.uk/staff/personal_pages/louisep/PEsymp/.

Organizing Committee

Louise Pryor (Chair), University of Edinburgh, louisep@aisb.ed.ac.uk; R. James Firby, University of Chicago, firby@cs.uchicago.edu; Steve Hanks, University of Washington, hanks@cs.washington.edu; Sam Steel, University of Essex, sam@essex.ac.uk.