

AAAI-98 Workshops

Reports of the Workshops held at the Fifteenth National Conference on Artificial Intelligence in Madison, Wisconsin

The Fifteenth National Conference on Artificial Intelligence (AAAI-98) was held in Madison, Wisconsin, on 26–30 July. The following four workshops were held in conjunction with the conference: (1) Case-Based Reasoning Integrations, (2) Learning for Text Categorization, (3) Predicting the Future: AI Approaches to Time-Series Problems, and (4) Software Tools for Developing Agents.

Case-Based Reasoning Integrations

The AAAI-98 Workshop on Case-Based Reasoning (CBR) Integrations, attended by approximately 45 people, focused on issues concerning multimodal reasoning systems that contain a CBR component. Although CBR is a general problem-solving architecture that is frequently integrated with other reasoning methods to solve complex tasks, no previous workshop attempted to characterize CBR integration issues.

Workshop highlights included invited talks by Mary Lou Maher (University of Sydney) on design, Janice Glasgow (Queens University) on computational imagery, and Marc Goodman (Continuum Software, Inc.) on World Wide Web applications of CBR integrations. These talks provided the audience with a breadth of suggestions on how CBR can assist, and benefit from, integrations with other reasoning methods for a variety of tasks.

We also reserved time for focused discussion periods on CBR integrations with rule-based reasoning, constraint-based reasoning, and planning that were led by experts in these

respective areas. These periods helped to summarize research in these areas and, in particular, to focus discussion on important open research issues. A final panel on the synergistic effects of CBR integrations examined what types of benefit can accrue from incorporating CBR into a multimodal rea-

The Fifteenth National Conference on Artificial Intelligence (AAAI-98) was held in Madison, Wisconsin, on 26–30 July. The following four workshops were held in conjunction with the conference: (1) Case-Based Reasoning Integrations, (2) Learning for Text Categorization, (3) Predicting the Future: AI Approaches to Time-Series Problems, and (4) Software Tools for Developing Agents.

soning system, what the significance of these synergies is, how they can be realized, and what their costs are. This panel and the other discussion periods helped to balance discussion time with 10 presentations on novel research and a large poster session. Plans are now being pursued toward writing a summary article, based on contributions to this workshop, that characterizes CBR integration issues on control, knowledge representation, learning, evaluation, and task-specific benefits.

—David W. Aha
Naval Research Laboratory
Jody J. Daniels
Lockheed Martin
Advanced Technology Laboratories

Learning for Text Categorization

The immense growth of the web has caused the amount of text available online to skyrocket. Seeking to help users better navigate and organize this information has led many researchers in a variety of communities to work on automated learning techniques for text. The AAAI-98 Workshop on Learning for Text Categorization brought together researchers from many of these fields, such as information retrieval, machine learning, Bayesian networks, and natural language processing, to share their different experiences in tackling similar problems.

Fernando Pereira (AT&T Labs-Research) began the workshop with an invited talk on the use of finite-state machines in text-processing problems. Then, discussions and presentations of various algorithmic issues in text categorization ensued, including work on new learning algorithms for text categorization and discussions of new problems that include dealing with changing user preferences, using active learning in text categorization, and using learning to better order retrieval results.

In another invited talk, Jaime Carbonell (Carnegie Mellon University) addressed the task of automatically detecting the emergence of new topics in an information stream. Subsequently, the workshop turned from algorithmic issues to representational issues. Specifically, several researchers made the point that making use of linguistic structure, as well as using stylistic and nontextual features of documents, can improve categorization performance.

During the workshop, two main thrusts for further investigation emerged. The first point dealt with using more expressive algorithms and representations to better capture linguistic structure. The second direction focused on gaining a better theoretical understanding (for example, bias-variance issues) of text categorization. Overall, the workshop was successful at shedding new light on current research findings and helping set directions for future work.

—Mehran Sahami
Epiphany Inc.

Predicting the Future: AI Approaches to Time-Series Problems

The Workshop on AI Approaches to Time-Series Problems, jointly sponsored by the Fifteenth National Conference on Artificial Intelligence (AAAI-98) and the International Conference on Machine Learning (ICML-98), was held in Madison, Wisconsin, on 27 July 1998. The organizing committee consisted of Andrea Danyluk of Williams College, and Tom Fawcett and Foster Provost, both of Bell Atlantic Science and Technology. There were approximately 30 attendees.

The goal of the workshop was to bring together AI researchers who study time-series problems along with practitioners and researchers from related fields. These problems are of particular interest because of the large number of high-profile applications today that include historical time series (for example, prediction of market trends, crisis monitoring). The focus was primarily on machine-learning and data-mining approaches, but perspectives on statistical time-series analysis and state-space analysis (for example, work on hidden Markov models) were also included. These communities arguably overlap significantly (and indeed, much work on state-space analysis, for example, has fallen under the heading of machine learning). Our goal was to make researchers and practitioners not only aware of approaches similar to their own but also aware of methods from other communities that might be applied effectively to their problems.

Time-series problems include segmentation and labeling of time series as well as prediction. Classical time-series prediction involves fitting a function to a numeric time series to predict future values (for example, the prediction of a stock's price given its past performance). Time-series prediction might also be required for problems involving categorical, rather than numeric, data. For example, one might be interested in predicting the next action of a computer user given a history of user actions.

In many situations, the problem is

not to predict future values of a time series but, rather, to label the series. For example, cardiologists examine electrocardiograms to diagnose arrhythmias. Problems of segmentation and labeling often occur together. The problem might be to (1) assign a label to an entire time series or (2) segment the series into subsequences and label them individually. The second problem is more difficult because there is temporal information to consider both within each subsequence itself and among the various subsequences. In the extreme case, the subsequences might be single atomic events, where temporal information exists among the events, but there is no temporal information within the event itself.

The labeling problems we described assume that there are fixed boundaries that define the pieces of the time series that are to be labeled. The boundaries, however, might not be fixed. For example, consider the problem of identifying and labeling the stages of a multistage flight plan given a pilot's command sequence. Here, the boundaries of the subsequences will not be fixed for all flight plans but will be variable. The problem becomes not only labeling the stages but also identifying the boundaries themselves.

In other cases, the identification of the boundary can be paramount. For example, consider the problem of detecting credit card fraud. Time-series information exists in the form of a stream of credit card transactions for an account, and the problem is to decide at any given time whether the account has been defrauded. In real time, this process would involve examining (sub)sequences of account activity and determining, as soon as possible, whether fraudulent activity exists. Here, the problem is not only to label accurately but also to identify as closely as possible the point where fraudulent behavior begins, so that use can be stopped and losses minimized.

The working notes contain 16 papers, 7 of which were selected for presentation at the workshop. In keeping with the goal of making the workshop a resource for researchers and practitioners in this area, the working notes included a small selection of rel-

evant papers that also appear elsewhere. The papers presented at the workshop span the range of problems described here.

In addition to presentations of papers in the proceedings, there were two invited talks. Leslie Kaelbling (Brown University) gave a brief tutorial on reinforcement learning and how it might be applied to problems in time series. Padhraic Smyth (University of California at Irvine) discussed the importance of not dismissing well-understood techniques from statistics, at least as a first method of attack when addressing new problems.

Discussions were held throughout the day, and a number of common issues arose. One recurring issue is the difficulty of obtaining data. Although time-series problems are becoming increasingly important and prominent, much of the data of interest (for example, financial or marketing data) are difficult to obtain because of their sensitive nature. User-profiling data are similarly difficult to share because of privacy concerns. A goal of the researchers at the workshop is to try to overcome some of these difficulties and establish a common repository of time-series data. By making these data commonly available, researchers can test new algorithms more easily, and the field can benefit from replication of experiments and comparison of results.

—Andrea Danyluk
Williams College

Tom Fawcett
Foster Provost
Bell Atlantic Science and
Technology

Software Tools for Developing Agents

Increasing interest in the design and use of intelligent agents has resulted in the development of a large number of software architectures and tools. However, to date, relatively little work has been done on theories or methodologies for the development of agent-based systems.

The aim of the AAAI-98 Workshop on Software Tools for Developing Agents was to relate generic issues in

Welcome to ijcai

July 31 – August 6 in Stockholm, Sweden

Sixteenth International Joint Conference on Artificial Intelligence

The Registration Brochure is now available

IJCAI-99 is composed of

- 30 Workshops July 31–August 2
- 20 Tutorials August 1–2
- Technical Sessions August 3–6
- Exhibition August 3–5
- RoboCup-99 July 27–August 6

These and additional affiliated conferences and workshops are described in the brochure.

Contact

Department of Computer and Systems Sciences,
Stockholm University and Royal Institute of Technology,
attn: IJCAI-99, Electrum 230, SE-164 40 Kista,
Sweden

phone: + 46 8 16 16 12

e-mail: IJCAI-99@dsv.su.se

<http://www.ijcai.org/ijcai-99>



IJCAI-99 is sponsored by IJCAI and the Scandinavian AI Societies

the design and development of intelligent agents to a discussion of existing and future software tools, especially integrated tool kits, with a view to providing a unified conceptual framework for analyzing requirements, assessing strengths and weaknesses of existing tools, and providing guidelines and suggestions for future tool kits. Although several conferences and workshops have dealt at some point with agent development, their wide scope has resulted in only limited discussion of tool kits and related issues, and this workshop was the first to focus on these aspects of agent development.

A broad range of work was presented, ranging from task and architecture neutral classifications of agent-based systems to implementation case studies and analyses of individual tool kits. Two main approaches to agent development were identified: (1) packages or libraries that extend an existing programming environment with addi-

tional facilities, such as sending and receiving messages in KQML, and (2) frameworks or shells that provide some form of *basic agent* that can be refined and extended for a particular problem. Inevitably, some work mixed these strategies, for example, one paper traced the migration of tools from user libraries to supported components within an agent-development framework. Other topics included agents programmed or “trained” directly by end users and the identification of agent-design patterns. With a few exceptions, most of this work is being done in mainstream languages such as Lisp, C++, or JAVA rather than special-purpose agent-development languages.

Discussion sessions focused on general issues affecting the design of tool kits and their performance rather than on specific applications or the embodiment of a particular agent theory. Major themes to emerge included the role of formal models of agents in the

development of tool kits and applications and the use of conversation protocols to structure communication between agents.

More information about the workshop, including copies of the workshop papers, can be found on the workshop web page: www.cs.bham.ac.uk/~bsl/aaai-98.

—Brian Logan

University of Birmingham

Jeremy Baxter

United Kingdom Defense

Evaluation and Research Agency