TUTORIAL PROGRAM OVERVIEW

Tutorials designated <<SA>> will be held Sunday, August 24, from 9:00am-1:00pm. <<SP>> tutorials will be held Sunday, August 24, from 2:00pm-6:00pm. <<MA>> tutorials will be held Monday, August 25, from 9:00am-1:00pm. <<MP>> tutorials will be held Monday, August 25, from 2:00pm-6:00pm

Cooperative Information Systems: Agents Meet Databases (SA1)
Munidar P. Singh and Michael N. Huhns

Modeling with Defaults: Causal and Temporal Reasoning (SA2)
Hector Geffner

Constraint Satisfaction Problems and Objects (SA3)
Pierre Roy and Francois Pachet

Logic and Learning (SA4)
Nada Lavrac and Luc De Raedt

Intelligent Multimedia Interface Agents (SP1)
Wolfgang Wahlster and Elisabeth Andre

Resource-Bounded Reasoning (SP2)
Shlomo Zilberstein

The CommonKADS Methodology (SP3)
Andre Valente and John Kingston

Robot Learning (SP4)
Sebastian Thrun

Practical Planning Systems (MA1)
Steve Chien and Brian Drabble

Qualitative Spatial Reasoning (MA2)
Tony Cohn

Genetic Programming (MA3)
John R. Koza and David Andre

Machine Learning and Information Filtering on the Internet (MA4)
Michael Pazzani

Economically Founded Multiagent System (MA5)
Tuomas Sandholm

Data Mining and Knowledge Discovery in Databases (MP1)
Usama Fayyad and Evangelos Simoudis

Pragmatics of Default Reasoning and Theory Change (MP2)
Grigoris Antoniou, Abhaya Nayak and Aditya Ghose

Neural Networks for Structured Knowledge (MP3)
Franz Kurfess and Alessandro Sperduti

Case-Based Reasoning (MP4)
David Leake and Ralph Barletta

TUTORIAL DESCRIPTIONS

SA1:
Cooperative Information Systems:
Agents Meet Databases

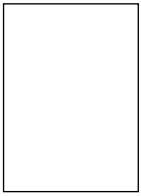

Speakers: Munidar P. Singh and Michael N. Huhns

Course Description:
Agents have been gathering an increasing amount of attention lately in the research community, from funding agencies, and even in the lay press. Successful agent applications, especially in information-rich domains, will depend not only on AI techniques but also on a solid understanding of the underlying database issues.

The science of Cooperative Information Systems (CIS) synthesizes AI and database results to develop effective systems - those composed of existing, heterogeneous components within an enterprise, as well as those whose components are independently developed and designed to behave autonomously in separate enterprises. Applications of CIS include telecommunications, virtual enterprises, logistics, healthcare, and manufacturing automation, to name but a few. CIS is distinguished from other AI work in agents by involving robust database techniques for capturing and using semantics through abstractions such as data models, ontologies, transactions, relaxed transactions, and workflows.

This tutorial will present the necessary concepts, architectures, theories, techniques, and infrastructure to build cooperative information systems. It will include a comprehensive overview of the state of the art in agent applications in distributed, heterogeneous databases.

We believe a tutorial such as this one, which combines AI, databases, and distributed computing, is essential for anyone trying to quickly come up to speed on a vast research area. This tutorial will guide practitioners by describing implemented, tested agent-based approaches to large-scale information access and management. It will introduce graduate students and other researchers to a new area with lots of exciting and important problems.

Prerequisite Knowledge: The tutorial is self-contained. No special background is assumed.

Munidar P. Singh (Ph.D., Texas, 1993) directs the Database Laboratory at North Carolina State University. He authored the book "Multiagent Systems" and several papers on agents and databases. Dr. Singh is the Americas program chair for the 1997 International Conference on Cooperative Information Systems, and the general chair for the 1997 International Workshop on Agent Theories, Architectures, and Languages. He serves on the editorial board for IEEE Internet Computing and as an editor of Agent-Based Computing. Dr. Singh received the US National Science Foundation CAREER Award.

Michael N. Huhns (Ph.D., South California, 1975) directs the Center for Information Technology at the University of South Carolina. He edited the books "Distributed Artificial Intelligence", volumes 1 and 2, and authored over 100 papers and reports. Dr. Huhns has served on numerous conference committees and advisory boards and is an Associate Editor for IEEE Expert and ACM Transactions on Information Systems. He serves on the editorial boards for the International Journal of Cooperative Information Systems, the Journal of Intelligent Manufacturing, and IEEE Internet Computing.

SA2:
Modeling with Defaults: Causal and Temporal Reasoning

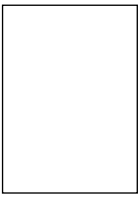

Speaker: Hector Geffner

Course Description:
A robot pushes a block and expects the block to move. The block however does not move. He pushes again but harder. The block moves.

Inferences of this type are easy for people but hard for robots. Part of the problem is that the modeling languages used in AI do not deal with uncertainty in a natural way. Logical languages for example do not handle uncertainty at all, while probabilistic languages deal with uncertainty at a precision and cost that is seldom needed.

Default languages are a new type of modeling languages that aim to fill the gap that exists between logical and probabilistic languages, providing modelers with the means to map soft inputs into soft outputs in a meaningful and principled way. Default models combine the convenience of logical languages, the flexibility and clarity of a probabilistic semantics, and the

transparency of argumentation algorithms. The goal of the tutorial is to provide a coherent and self-contained survey of such work.

We view default reasoning in two ways: as an extended form of deductive inference and as a qualitative form of probabilistic inference. In each case, we lay out the main concepts, intuitions and algorithms. We then consider the specific problems that arise when reasoning about causality and time, and analyze what works, what doesn't work, and why. We make use of the basic ideas that underlie two probabilistic models: Bayesian Networks and Markov Processes. This allows us to shed light on a number of issues like the distinction between laws and facts, the role of causality, and the conditions for efficient reasoning.

We also illustrate the use of default languages for modeling in areas such as qualitative reasoning, decision making, and planning and control.

Prerequisite Knowledge: The tutorial is intended for people interested in common sense modeling, planning, decision and control. There are no prerequisites except a basic knowledge of logic and probabilities.

Hector Geffner got his Ph.D. in UCLA with a dissertation on Default Reasoning that was co-winner of the 1990 ACM Dissertation Award. Then he worked as Staff Research Member at the IBM T.J. Watson Research Center in NY for two years before returning to the Universidad Simon Bolivar, in Caracas, Venezuela where he currently teaches. Hector Geffner has served in the program committee of the major AI conferences and is a member of the editorial board of the Journal of Artificial Intelligence Research.

SA3:
Constraint Satisfaction Problems and Objects

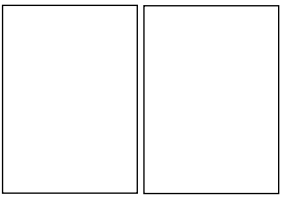

Speakers: Peirre Roy and Francois Pachet

Course Description:
Combinatorial optimization is a powerful paradigm to solve complex problems. It has a wide range of applications such as planning, scheduling, resource sharing, in a wide range of domains (transportation, production, mass marketing, network optimization, human resources management). CSP techniques provide efficient algorithms to prune search spaces. On the other

hand, object-orientation is an acknowledged paradigm to represent richly structured domains. The combination of the two paradigms yields powerful programming languages. This tutorial aims at introducing the area of CSP and their use in conjunction with an object-oriented programming language. The tutorial will describe the range of applications, and give a comprehensive overview of technical issues.

The integration of CSP and objects raises two main issues. First, it should provide an efficient implementation of existing algorithms to solve complex problems. We illustrate this point with the design of a constraint solver, called BackTalk, a canonical integration of CSP with Smalltalk based on the systematic reification of the main concepts of constraint satisfaction: domains, variables, constraints, problems and algorithms. This design allows to integrate specialized algorithms for each constraint.

The second point concerns the representation of knowledge to speed up the search. In the context of embedded object-oriented CSPs, this aspect has two sides. First, we show how the design of a CSP involving objects can drastically influence the performance of the resulting system. This point is illustrated with a system that performs musical harmonization of melodies. Second, we show how domain specific knowledge can be expressed and exploited by the standard CSP mechanism to avoid exploring useless branches of the search space. This point will be illustrated by a system that generates crossword grids.

Prerequisite Knowledge: Attendees should be familiar with object-oriented programming, but not necessarily with constraint programming.

Pierre Roy is pursuing a Ph.D. thesis at the Laforia laboratory, under the supervision of Francois Pachet. He is the main author of the BackTalk system that will be used throughout the tutorial for demonstrations. He has realized an automatic harmonization system built with BackTalk and has published several papers dealing with the integration of constraint satisfaction techniques with objects. He has presented a tutorial on these matters at the conference Expert Systems'1995 in Cambridge (UK), and at the European Smalltalk Users Group (ESUG) in Lausanne (Switzerland) in 1996.

Francois Pachet (Ph.D., Eng.) is associate Professor at University of Paris 6, Laforia-IBP. He is specialized in knowledge representation using object-oriented techniques. He is the author of several papers describing the integration of artificial intelligence techniques in object-oriented languages. He presented various tutorials and organized several workshops at the OOPSLA

conference on this theme (embedded object-oriented production systems, metamodeling in 95).

SA4:
Logic and Learning

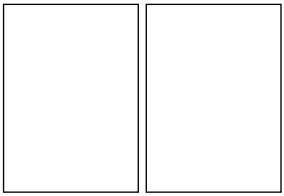

Speakers: Nada Lavrac and Luc De Raedt

Course Description:
Since the very start of machine learning, logic has been very popular as a representation language for inductive concept-learning and the possibilities for learning in a first order representation have been investigated. This is due to the fact that first order logic extends propositional representations and therefore also the scope of machine learning. In database terminology, propositional techniques learn from a single relation in a relational database, whereas first order approaches cope with multiple relations. Learning in first-order logic is therefore of special interest to researchers and practitioners in machine learning, data mining, knowledge discovery in databases, knowledge representation and logic programming.

The tutorial will provide a complete survey of "Logic and Learning" and will concentrate on learning in first order logic. Early research involved structural matching, least general generalization, model inference, and theory restructuring. Recently, the area of logic and learning has concentrated on Inductive Logic Programming, which studies inductive machine learning within the representation of logic programming.

The course gives an overview of the history, techniques and applications of logic and learning. Logical and algorithmic foundations of this field are emphasized, and illustrated by means of well-known systems and algorithms. This includes: learning as search, logical representations, operators and settings for induction, state-of-the-art methods and systems, and applications in knowledge discovery and logic programming. For use in data mining and knowledge discovery, pointers to the available public domain systems are provided and discussed.

Prerequisite Knowledge: The course assumes basic knowledge of artificial intelligence and logic (e.g. some notions about Prolog and/or the predicate calculus).

Nada Lavrac is a research associate at the J. Stefan Institute, Ljubljana (Slovenia). Her main research interest is in inductive

logic programming and medical applications of machine learning. She is co-author and editor of several books published by Sigma Press, MIT Press, Kluwer and Springer, including "Inductive Logic Programming: Techniques and Applications" (Ellis Horwood 1994). She was a coordinator of ILPNET, the European Scientific Network in Inductive Logic Programming (1993-96).

Luc De Raedt is a post-doctoral researcher and a (part-time) assistant professor at the Katholieke Universiteit Leuven (Belgium). His main interest is in inductive logic programming. He is a coordinator of the European ESPRIT III and IV projects on Inductive Logic Programming, was chair of the ILP-95 workshop, and he has published two books on inductive logic programming. He has given tutorials and invited talks on ILP at ISMIS-93, ILPS-93, MSL-96, LOPSTR 96, etc.

SP1:
Intelligent Multimedia Interface Agents

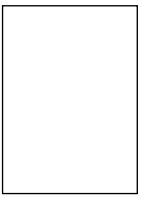

Speakers: Wolfgang Wahlster and Elisabeth Andre

Course Description:
The goal of this tutorial is to survey a new generation of highly personalized agent-based or assistant-like user interfaces. Personalization refers to the ability of an interface to adapt its behavior to the information needs, interaction styles, and media preferences of individual users in particular situations. In the context of the World-Wide Web it wouldn't make sense to delegate the task of personalizing presentations to the information providers. One reason is that information providers would have to anticipate all possible users and situations in order to provide adequate presentation formats.

In the tutorial, we will discuss recent developments in the design of intelligent multimedia interfaces that go beyond the standard canned text, predesigned graphics and prerecorded images and sounds typically found in commercial multimedia systems of today. We will show that it is possible to adapt many of the fundamental concepts developed to date in NL processing, user modeling and discourse research in such a way that they become useful for multimedia presentations as well. We will address key applications such as communication assistants for the Internet, multimedia helpware, information retrieval and analysis, authoring, training, monitoring, and decision support. There is also a peripheral aspect of personalizing user interfaces when

the system personifies itself audio-visually, e.g. as an animated life-like character.

The tutorial will introduce the technology for the development of animated interface agents who play the role of communication assistants explaining, commenting and highlighting the material to be presented. The tutorial will be augmented by numerous videos and interactive demos of research prototypes and commercial applications.

Prerequisite Knowledge: The tutorial assumes no prior knowledge on multimedia interface agents, but a basic knowledge of AI concepts, including AI planning, NL processing, graphics generationand user modeling, will enhance the value of this course for participants.

Wolfgang Wahlster is the Director of the German Research Center for AI (DFKI) and a Professor of Computer Science at the University of Saarbruecken. He is a AAAI Fellow, a Trustee of IJCAII, the Chair of ECCAI and the Programme Co-Chair of ACL/EACL-97. He has published more than 150 technical papers on intelligent user interfaces. His current research includes multimodal interfaces, user modeling, and speech translation. He is a Co-Editor of the Readings in Intelligent User Interfaces.

Elisabeth Andre has been a member of the `Intelligent User Interfaces' group at DFKI since 1989. Her research interests include: multimedia communication, intelligent user interfaces, natural-language processing and life-like characters. Since January 1994, she has been member of the executive board of the ACL Special Interest Group on Multimedia Language Processing (SIGMEDIA).

Both presenters are actively involved in numerous industrial projects dealing with various applications of intelligent user interface technology.

SP2:
Resource-Bounded Reasoning

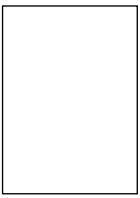

Speaker: Shlomo Zilberstein

Course Description:
Resource-bounded reasoning is an emerging field within artificial intelligence that addresses one of its primary challenges: how to embed complex reasoning components in real-world

applications. The need to employ resource-bounded reasoning techniques is based on a simple, but general, observation. In many situations, the computational resources required to reach an optimal decision reduce the overall utility of the result. This observation covers a wide range of applications such as automated diagnosis and treatment, signal interpretation, combinatorial optimization, probabilistic inference, mobile robot navigation, visual tracking, graphics, and information gathering. What is common to all these problems is that it is not feasible (computationally) or desirable (economically) to compute the optimal answer.

Moreover, taking the cost of computation into account is not an easy task, since the "optimal" level of deliberation varies from situation to situation.

A multitude of resource-bounded reasoning techniques have been developed in recent years exploiting new computational techniques that allow small quantities of computational commodities-such as time, memory, or information-to be traded for gains in the value of computed results.

This tutorial will examine the benefits and limitations of recently developed techniques including anytime algorithms, flexible computation, memory-bounded search, imprecise computation, and design-to-time scheduling. Topics which will be covered include: introduction and historical background, types of computational tradeoffs in reasoning and search, representation and measurement of computational tradeoffs, embedding flexible computation components in large systems, run-time assessment and prediction of solution quality, monitoring and control of computational resources, performance evaluation of resource-bounded reasoning systems, a brief survey of successful applications, and current research directions.

Prerequisite Knowledge: This tutorial is designed for both researchers interested in fundamental issues in resource-bounded reasoning and practitioners with a primary interest in applications. The tutorial will be self-contained and requires basic familiarity with AI (search and automated reasoning) and probability theory.

Shlomo Zilberstein is an Assistant Professor of Computer Science and the head of the Resource-Bounded Reasoning Research Group at the University of Massachusetts. He received his BA in Computer Science (1981) from Israel Institute of Technology summa cum laude, and his Ph.D. in Computer Science (1993) from the University of California at Berkeley. Prof. Zilberstein has 15 years of experience in research and

development of real-time intelligent systems. He has published numerous articles on resource-bounded reasoning and has presented his work at major conferences. He has also organized several successful workshops in this area including the IJCAI-95 Workshop on Anytime Algorithms and Deliberation Scheduling, the 1996 AAAI Fall Symposium on Flexible Computation, and the AAAI-97 Workshop on Building Resource-Bounded Reasoning Systems.

SP3:
The CommonKADS Methodology

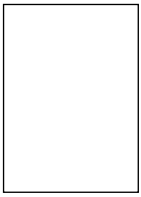

Speakers: Andre Valente and John Kingston

Course Description:
When knowledge is acquired for the construction of a knowledge based system (KBS), it must be represented and structured in such a way that it can be used to analyse existing structures and approaches, and to produce a coherent design for a knowledge based system. The separation of analysis and design is a key element in both accurate modelling of expert knowledge in an implementation-independent manner (allowing the possibility of reusing knowledge models between applications) and in making well-justified decisions about the best implementation approach for a particular problem. The CommonKADS methodology provides a principled approach to representing, organising and transforming acquired knowledge; in other words, it is an *enabling technology* to support and guide the construction of knowledge based systems.

The goal of the tutorial is to introduce the audience to a disciplined approach to developing knowledge based systems based on the CommonKADS methodology. The course is intended for knowledge engineers and other technical specialists interested in methods for developing knowledge based systems. On completion, the audience will understand the benefits of modelling knowledge as an intermediate step between knowledge acquisition and implementation; know how to produce a set of knowledge models by following the CommonKADS methodology; be able to identify the task types of potential and actual KBS applications; and understand the basis of good design decisions.

Prerequisite Knowledge: An awareness of popular knowledge representation and inferencing techniques is expected. Some knowledge of KBS development tools and a small amount of programming experience is also desirable.

Andre Valente is a computer scientist at the USC Information Sciences Institute. Prior to his work at ISI he participated, at the University of Amsterdam, in the design of the Common KADS methodology and the CommonKADS Library in particular. He has also worked for major industrial corporations, performing applied research and development in knowledge engineering. He holds a Ph.D. (Computer Science) from the University of Amsterdam (1995). His research interests are knowledge engineering, knowledge acquisition, and planning.

John Kingston is a Senior Computer Scientist within the Knowledge Engineering Methods Group at the Artificial Intelligence Applications Institute (AIAI), University of Edinburgh. He has developed several commercial knowledge based systems, presented a range of commercial training courses, and provided consultancy on knowledge based systems in the UK, Europe and the USA. He also publishes frequently, and was awarded first prize for Best Technical Paper at the BCS Expert Systems '93 conference for a paper on applying the CommonKADS methodology.

SP4:
Robot Learning

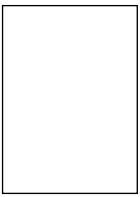

Speaker: Sebastian Thrun

Course Description:
Robot learning is concerned with algorithms that enable robots to improve their performance with experience. Learning can make up for lack of knowledge when programming robots, such as lack of exact task specifications, lack of environment models, lack of sensor models, or lack of effective control strategies. Traditionally, the field of robotics has paid little attention to robot learning. With a new generation of intelligent service robots in close reach, and with the number of success stories increasing, learning is likely to become a fundamental part of robotics.

This tutorial will provide an introduction into the basic algorithms and techniques used in robot learning. It will cover recent work on learning models, learning control, and probabilistic reasoning. In addition, it will highlight some recent success-stories of robot learning and give some guidance for applying robot learning in practice.

The tutorial is targeted towards students, engineers, scientists,

and teachers who are new to the field of robot learning, but who would like to get an overview of the field (and who would like to share my excitement).

Prerequisite Knowledge: Basic knowledge in robotics, machine learning, or statistics will be helpful but is not required.

Sebastian Thrun is a research faculty member at Carnegie Mellon University. His research interests lie in the areas of machine learning, neural networks and robotics. Thrun received his Ph.D. in 1995 and his M.Sc. in 1993, both from the University of Bonn in Germany. He is a consultant for several companies, including Real World Interface Inc., a leading US mobile robot manufacturer. Thrun recently co-edited a special issue of the journal Machine Learning and a book on "Robot Learning," and is currently editing another book entitled "Learning to learn".

MA1:
Practical Planning Systems

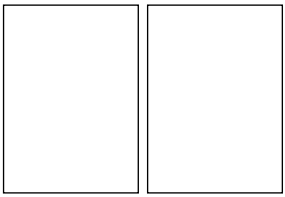

Speakers: Steve Chien and Brian Drabble

Course Description:
Automated planning is the generation of a low-level sequence of actions to achieve some desired world state while obeying domain constraints. Planning systems can be used to automate procedure generation problems and have been applied to such diverse fields as science data analysis, image processing, crisis response, space payload operations, and operating a network of communications antennas. Automated planning technology has the potential to reduce operations costs, decrease manual errors, and reduce dependency on key personnel.

This tutorial will cover the basic concepts in domain-independent artificial intelligence planning including: search, representing planning knowledge, plan and state space planning, operator-based planning and hierarchical task network planning. Advanced concepts such as planning and scheduling, decision theoretic planning, and mixed initiative planning will also be briefly discussed.

Important questions relevant to planning will be covered in the tutorial such as:
- Are planning techniques applicable to my problem?
- If so, what are the most appropriate planning representations and techniques to use?

- How to acquire, verify, and maintain, my planning knowledge base?
- How to embed a planning system into an operational setting?

The tutorial will be broadly beneficial to a number of different groups, in particular:
- AI practitioners seeking a thorough overview of the state-of the art in AI planning technology and key issues in fielding applications.
- Planning and related AI researchers seeking an overview of the current state of the art in AI planning and insights into key bottlenecks in fielding AI planning systems.

Prerequisite Knowledge: Knowledge of basic concepts from Artificial Intelligence will be presumed : search, expert systems, logic-like representations. Familiarity with some planning and scheduling systems, basic search strategies, reactive systems, and/or scripting languages would be helpful but not essential.

Steve Chien is Technical Group Supervisor of the Artificial Intelligence Group, at the Jet Propulsion Laboratory, California Institute of Technology where he leads efforts in automated planning and scheduling for spacecraft mission planning, maintenance of space transportation systems, and Deep Space Network Antenna operations. He holds a B.S., M.S., and Ph.D. in Computer Science, all from the University of Illinois. Dr. Chien is also an Adjunct Assistant Professor in Computer Science at the University of Southern California.

Brian Drabble is a Research Associate at the Computational Intelligence Research Laboratory (CIRL) at the University of Oregon. His current responsibilities are to transition the planning and scheduling research being undertaken at CIRL into industry, commerce and military applications. Previous to joining CIRL he spent 8 years as a member of the Artificial Intelligence Applications Institute at the University of Edinburgh. His responsibilities included being project leader and co-principal investigator on the O-Plan project which is part of the $66 million DARPA/Rome Laboratory Planning and Scheduling Initiative. In addition he has worked with a number of clients including Toshiba, Hitachi, European Space Agency, and the British Government, to bring intelligent planning and scheduling into their organisations and products. Brian Drabble holds a B.Sc. from Staffordshire University and a Ph.D. from Aston University both in the U.K

MA2:
Qualitative Spatial Reasoning

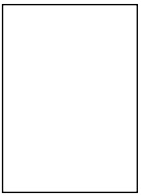

Speaker: Tony Cohn

Course Description:
While the field of Qualitative Reasoning is now well established it is only recently that there has been substantial investigation of calculi suitable for representing and reasoning about space in a qualitative way. This tutorial will survey the state of the art in qualitative spatial representation and reasoning techniques, including mechanisms for reasoning about spatial change. We will look at how these ideas may be applied in a wide variety of domains, from robotics and high level vision to the semantics of natural language expressions and of visual programming languages though to Geographical Information Systems. The tutorial will conclude with a look at open research issues.

The aim of the tutorial is to present the state of art in Qualitative Spatial Reasoning to the non specialist; this could either provide a suitable introduction for the intending research student, or for a practitioner in another research area or application domain who may be interested in exploiting QSR techniques. The presentation is expected to include practical demonstrations of techniques and applications including video footage. There will be ample opportunity for tutees to ask questions.

After an introduction covering basic ontology, some history and motivation, the tutorial syllabus will cover major aspects of qualitative spatial representation (including topology, distance and size, orientation and shape), uncertainty and vagueness, spatial change, reasoning techniques, applications and research issues.

Prerequisite Knowledge: Some basic knowledge of logic would be helpful but not essential for understanding the tutorial.

Tony Cohn is Professor of Automated Reasoning at the University of Leeds and has been leading a research group on qualitative spatial reasoning since the late 1980s. He coordinates the European network SPACENET on qualitative spatial reasoning and has given a number of invited tutorials and lectures on the subject in various international AI and Geographical Information Systems forums. He has been chairman of the European and British AI societies and programme chair of several AI conferences.

MA3:
Genetic Programming

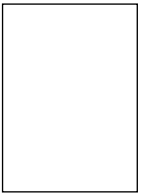

Speakers: John R. Koza and David Andre

Course Description:
This tutorial will introduce participants to the ideas and applications of genetic programming (GP) -- an automated search procedure based on the mechanics of natural genetics and natural selection. Genetic programming is a domain-independent technique for automatic programming that evolves computer programs that solve, or approximately solve, problems. Starting with a primordial ooze of thousands of randomly created programs composed of functions and terminals appropriate to a problem, a population of computer programs is progressively evolved over many generations by applying the Darwinian principle of survival of the fittest, a sexual recombination operation, and occasional mutation.

Genetic programming has found successful applications in a wide variety of different areas of artificial intelligence, including engineering design, planning, robotics, programming of independent agents, distributed artificial intelligence, modeling, system identification, forecasting, empirical discovery, data mining, optimal control, pattern recognition, game theory, optimization, structural design, molecular biology, creation of mental models, and knowledge reuse.

We will briefly review the properties and mechanics of genetic programming, and then discuss the techniques and methods that have been employed in recent years to apply genetic programming to a variety of difficult real-world problems. Topics include hierarchical, multi-part programs, automatically defined functions, the use of iteration, recursion, memory structures, mental models, architecture-altering operations, cellular encoding, implementation on parallel computers, genetic design of electrical circuits, genetically evolved assembly code, evolvable hardware, promising application areas for genetic programming, and directions for future research.

Prerequisite Knowledge : No knowledge about genetic algorithms, genetic programming or biology is required; however, a general familiarity with computers and programming is assumed. The basics of genetic programming will be provided and the tutorial will concentrate on intermediate and advanced topics.

John R. Koza is a Consulting Professor of Computer Science at Stanford University. He is author of two books on genetic programming: Genetic Programming: On the Programming of Computer by Means of Natural Selection (MIT Press, 1992) and Genetic Programming II: Automatic Discovery of Reusable Programs (MIT Press, 1994).

David Andre is currently doing research on genetic programming and artificial intelligence at UC Berkeley. He has published more than 20 papers on genetic programming and is working on an upcoming book. He has been researching genetic programming for the past five years.

MA4:
Machine Learning and Information Filtering on the Internet

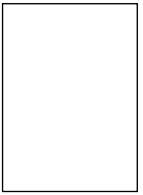

Speaker: Michael Pazzani

Course Description:
The vast amount of information available on the Internet has given rise to a number of agents for locating relevant, useful or interesting information for an individual. Such agents perform tasks such as prioritizing, filtering, or sorting electronic mail; filtering news group articles and locating interesting articles in unread newsgroups; guiding a user to find relevant information on the World Wide Web; notifying a user when a significant change occurs to a web site or providing access to information relevant to a user's current tasks.

To perform such tasks, a profile of the user's interests must be created. In this tutorial, we will focus on the learning and representation of user profiles, the methods for collecting user feedback, and the representation of information sources. This tutorial will review a variety the findings from several decades of research on information retrieval focusing on approaches to information filtering and classification. Next, machine learning approaches to classification will be described including decision trees, nearest neighbor algorithms, Bayesian classifiers and neural networks. We will discuss how they may be used to learn user profiles. The relationship between machine learning and classic approaches from information retrieval will be discussed. Finally, recent developments such as collaborative filtering, efficient rule learners, combining multiple models, weighted majority algorithms and infinite attribute models will be described.

The technology will be illustrated with examples from a variety of information agents including LIRA, NewsWeeder, WebWatcher, WebDoggie, InfoFinder, Inquery, Letizia, firefly, InfoFinder, Syskill & Webert, DICA and the Remembrance Agent.

Prerequisite Knowledge: The intended audience of this tutorial is practitioners and researchers interested in issues involved with applying machine learning and information retrieval algorithms to classification and ranking of information on the Internet. A familiarity with basic knowledge of mathematics and probability will be assumed.

Michael Pazzani received an M.S. degree in computer science specializing in Natural Language Processing in 1980, and a Ph.D. in computer science specializing in Machine Learning from UCLA in 1987. He is now a professor and department chair of Information and Computer Science at the University of California, Irvine. He has been active in Machine Learning research for the past decade with numerous publications in the IJCAI, AAAI, Cognitive Science and the International Machine Learning Conferences.

MA5:
Economically Founded Multiagent System

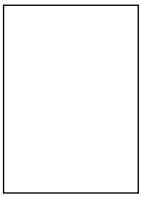

Speaker: Tuomas Sandholm

Course Description:
Multiagent systems research, a subfield of AI, studies the interactions of computational agents. These agents can represent different real world parties, and they can have different preference structures. Important applications include manufacturing planning and scheduling among multiple agile enterprises, markets for electricity, allocating and pricing bandwidth in multi-provider multi-consumer computer networks, network management, multiagent information gathering on the web, distributed vehicle routing among independent dispatch centers, electronic commerce, resource allocation in distributed operating systems, meeting scheduling, scheduling of patient treatments across hospitals, classroom scheduling, and planning and scheduling of multi-contractor software projects, to name just a few. Multiagent systems can save users' time, but they may also achieve better solutions (e.g. by enhanced negotiation and coalition formation)

than human agents can in combinatorially and strategically complex domains.

A key research goal is to design open distributed systems in a principled way that leads to globally desirable outcomes even though every participating agent only considers its own good and may act insincerely. The tutorial covers relevant results in AI, game theory, market mechanisms, voting, auctions, coalition formation, and contract nets. Emphasis is given to rigorous concepts, results, and algorithms - both classic ones from microeconomics and recent ones from the distributed AI community - that have direct applications to computational multiagent systems. Implementation experiences will be shared. Effects of different computational limitations (agents' bounded rationality) are discussed as a key feature that has not received adequate attention. Examples of real-world applications will be presented.

Prerequisite Knowledge: The tutorial equips the serious systems builder with rigorous techniques for making multiple self-interested agents cooperate efficiently. It also serves to familiarize newcomers and executive level participants with the issues involved. No prior knowledge is assumed in economics or multiagent systems. A general familiarity with computer science will be helpful.

Tuomas Sandholm is assistant professor of computer science at Washington University. He received the M.S. (B.S. included) with distinction in Industrial Engineering and Management Science from the Helsinki University of Technology, Finland, in 1991. From 1988 to 1992, he worked as a research scientist in the software industry. He earned the M.S. and Ph.D. degrees in computer science from the University of Massachusetts at Amherst in 1994 and 1996 respectively. He has published 23 refereed articles in AI Journal, IJCAI, AAAI and other forums - as well as numerous book chapters, technical reports and other papers. He has been a program committee member for five major conferences, and a reviewer for seven journals and numerous conferences. He has six years of experience designing efficient multiagent systems. This work has focused both on theory and implementations. He has also been involved in developing two fielded AI systems: a pension law expert system and a large-scale transportation optimization application.

MP1:
Data Mining and Knowledge Discovery in Databases

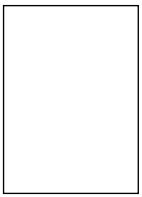

Speakers: Usama Fayyad and Evangelos Simoudis

Course Description:
Knowledge Discovery in Databases (KDD) is a rapidly growing AI field that combines techniques from machine learning, pattern recognition, statistics, databases, and visualization to automatically extract knowledge (or information) from lower level data (databases). This knowledge is subsequently used to support human decision-making, e.g., prediction and classification tasks, summarize the contents of databases, or explain observed phenomena. The use of KDD systems enables decision makers to automatically analyze the large and complex data sets collected today without requiring detailed prior knowledge about the data. Successful KDD systems have been implemented and are currently in use in financial modeling, fraud detection, market data analysis, astronomy, diagnosis, manufacturing, and biology.

This tutorial presents a comprehensive picture of current research paradigms in the field of KDD and examples from the state of practice. The tutorial provides an introduction to KDD, defines the basic terms and the relation between data mining and the KDD process, presents methods for data preparation and preprocessing, describes major data mining techniques from the fields of AI, pattern recognition, databases, and visualization, discusses major KDD systems from academia and industry, and provides a guide for developing a KDD system. In the process, the tutorial addresses such issues as role played by the various steps in the KDD process, e.g., sampling, data selection, projection and dimensionality reduction, extraction of patterns and models, and the use of extracted knowledge in decision-making.

Prerequisite Knowledge: There are no pre-requisites for this tutorial other than familiarity with basic concepts in AI.

Usama Fayyad is a Senior Researcher at Microsoft Research, a Distinguished Visiting Scientist at the Jet Propulsion Laboratory, Caltech, and an adjunct professor of computer science at University of Southern California. Prior to joining Microsoft he headed the Machine Learning Systems Group at JPL. He

received his Ph.D. in Computer Science (1991) from the University of Michigan, Ann Arbor. He was program cochairman of KDD-94 and KDD-95, general chair of KDD-96, and Editor-in-Chief of the Journal of Knowledge Discovery and Data Mining.

Evangelos Simoudis is Vice President of Decision Support Solutions at IBM and an adjunct professor of computer engineering at the Santa Clara University. Prior to joining IBM, Dr. Simoudis led the development and market introduction of the Recon data mining system, and led research on knowledge discovery in databases, and machine learning. Dr. Simoudis received his Ph.D. in Computer Science from Brandeis University. He is Editor-in-Chief of the Artificial Intelligence Review, and has served as Program cochairman of KDD-96.

MP2:
Pragmatics of Default Reasoning and Theory Change

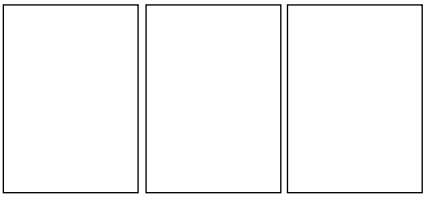

Speakers:
Grigoris Antoniou, Abhaya Nayak and Aditya Ghose

Course Description:
It is becoming increasingly apparent that the design of robust information systems must rely on principled methods for handling incomplete and changing information. As a result, there has been a resurgence of interest in the areas of nonmonotonic reasoning and theory change, focussing less on purely theoretical insights, and more on their ability to provide a starting point for the design of the next generation of tools for developing and managing complex information systems. This tutorial will explain what these core technologies are, why they are central to the development of real-world information systems and how they can be effectively deployed.

The tutorial will first present the formal principles underlying default reasoning, an approach to nonmonotonic reasoning which has been widely applied, and the logic of theory change. Using this formal background, the tutorial will examine applications of these techniques in a variety of domains, including diagnostic systems, intelligent scheduling systems, information retrieval and software engineering (requirements engineering, software maintenance). The tutorial will involve simple hands-on demonstrations using implemented systems for default reasoning and theory change.

Two sets of recent developments make this tutorial particularly timely. First, a number of recent studies have shown that a unified view of theory change and default reasoning leading to efficient implementations is possible. Second, an increasing number of useful applications of these technologies are being reported in the literature that build on these formal results. A clear understanding of the what, why and how of these issues is thus crucial for both researchers and practitioners interested in intelligent real-world information systems.

Prerequisite Knowledge: No prior knowledge is needed other than a basic understanding of classical logic.

Grigoris Antoniou is Senior Lecturer in Computing at Griffith University. His research interests include the logical foundations of computer science and artificial intelligence, and in particular nonmonotonic reasoning. He is co-author of "Logic: A Foundation for Computer Science", Addison-Wesley 1991, and author of a forthcoming book on Nonmonotonic Reasoning (MIT Press, 1997). He has published over 50 refereed papers at conferences and journals (including IJCAI, AAAI, Annals of Mathematics & Artificial Intelligence, Artificial Intelligence Review, Journal of Automated Reasoning).

Abhaya Nayak is a Senior Research Fellow at the Knowledge Systems Group, University of New South Wales. His research interests include belief dynamics, commonsense reasoning and counterfactual reasoning. He has published over a dozen papers in reputed journals (e.g. Journal of Philosophical Logic, Erkenntnis, Synthese) and conference proceedings. He holds a Ph.D. from the University of Rochester, U.S.A. and has held short-term teaching appointments at the University of Rochester and the State University of New York.

Aditya Ghose is Lecturer in Business Information Systems at the University of Wollongong, Australia. His research interests include default reasoning, theory change and constraint solving, and their applications in software engineering, planning and induction. He holds a Ph.D. in Computing Science from the University of Alberta, Canada. He has held research appointments with the Knowledge Systems Group at the University of Sydney and with the School of Computing at Griffith University, Brisbane and has been been a visiting researcher at the University of Tokyo.

MP3:
Neural Networks for Structured Knowledge

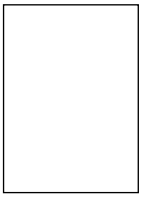

Speakers: Franz Kurfess and Alessandro Sperduti

Course Description:
This tutorial presents an overview of recent developments in using neural networks for representation and processing of structured knowledge. To date, mostly symbol-oriented methods have been used for this purpose; these methods, however, have problems with inexact and noisy data, inconsistent knowledge, brittleness, knowledge acquisition, and real-time constraints. Neural networks, on the other hand, are universal approximators, can perform automatic inference (learning), possess very good classification capabilities, and can deal with noise and incomplete data; they can frequently also be used as anytime-methods, where a (possibly non-optimal) result is available anytime during the evaluation, not only at the end.

In recent years, substantial progress has been made towards the use of neural networks for structured knowledge. Structured domains are characterized by complex patterns usually represented as lists, trees, and graphs of variable sizes and complexity. The basic problem is to overcome the limitation imposed by the fixed input size of a network when faced with the task of representing a graph of variable size and with an internal structure. Standard neural networks are well suited for dealing with unstructured patterns, and recurrent neural networks can be used to process sequences; a generalization of a recurrent neuron, the generalized recursive neuron, is capable of representing, classifying, and storing structured information very naturally.

This tutorial will present various approaches for the use of neural networks to deal with structured knowledge.

Prerequisite Knowledge: Attendees should be familiar with basic concepts of neural networks, or have background knowledge in a related area such as machine learning or statistics.

Franz Kurfess is co-director of the Software Engineering Lab and associate director of the Electronic Enterprise Engineering program at the Computer and Information Systems Department,

New Jersey Institute of Technology. His main area of research is the integration of systems based on various computation methods, e.g. symbolic and connectionist systems, and in particular logic and reasoning with neural networks.

Alessandro Sperduti is an Assistant Professor at the Computer Science Department, University of Pisa, Italy. His research area is neural networks, especially for representing structured information. He has written around 40 refereed papers mainly in the areas of neural networks, fuzzy systems, and image processing.

MP4:
Case-Based Reasoning

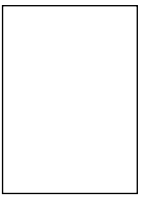

Speakers: David Leake and Ralph Barletta

Course Description:
Case-based reasoning (CBR) has become an established part of artificial intelligence, both as a means for addressing fundamental AI problems and as a basis for fielded AI technology. This tutorial defines the fundamental principles of CBR and shows how to apply them to real-world problems.

We start by presenting an overview of the case-based reasoning process, comparing and contrasting CBR to other AI methods to clarify motivations for the case-based approach. We then highlight key issues for CBR systems, principles for addressing them, and their role in high-impact applications and research systems. This discussion will be based on case studies of important CBR systems that reason and learn in multiple taskdomains. For each case study we talk about fundamental CBR issues the system addresses, consider how effective the system is in accomplishing its stated goal, and finally assess the contributions of the system to the state of the art in CBR.

Following this overview, we look at the major issues in designing and maintaining real-world CBR applications. We will define a methodology for building successful CBR systems in the real world, addressing tasks ranging from data collection and knowledge engineering to project scoping and user interface design. The tutorial closes by highlighting lessons learned from CBR research and applications, current challenges for case-based reasoning methods, how the challenges are being addressed, and opportunities for the next generation of CBR systems and their future impact.

The tutorial will clarify key principles of CBR, will supply applications developers with the information they need to begin selecting and applying CBR tools, and will provide experienced professionals with an in-depth view of key issues, methods, and emerging opportunities for applying case-based reasoning technology.

Prerequisite Knowledge: The tutorial is aimed at both researchers and application builders who are interested in how to integrate CBR into the systems they are building. Prior experience with CBR will certainly enhance the value of the tutorial but is definitely not required.

David Leake is an Assistant Professor of Computer Science at Indiana University. His primary research area is case-based reasoning, in which he has over 50 publications including the book "Case-Based Reasoning: Experiences, Lessons, and Future Directions" (AAAI Press, 1996). He chaired the AAAI-93 Workshop on Case-Based Reasoning, presented the opening tutorial at the First International Conference on Case-Based Reasoning in 1995, and is co-chair of the 1997 Second International Conference on Case-Based Reasoning. He received his Ph.D. in Computer Science from Yale University.

Ralph Barletta is Chief Scientist at Inference Corporation, a leading vendor of software tools for problem resolution using CBR technology, where he defines and oversees research and development for CBR products. He was President of Case Data Solutions Inc., a consulting firm specializing in AI and CBR, and oversaw development of ReMind, one of the first commercial CBR tools. His numerous publications include the book "A Review Of Industrial CBR Tools," which he co-authored in 1995. He presented the opening tutorial at the Second European Workshop on Case-Based Reasoning in 1994. He received his B.A. in Business Administration and M.S. in Computer Science from Rutgers University.