

Abstracts of Recent PhDs

An Intelligent Broker Architecture for Pervasive Context-Aware Systems

Candidate: Harry Lik Chen

Institution: Department of Computer Science and Electrical Engineering University of Maryland, Baltimore County, MA, USA

Supervisor: Timothy W. Finin

Year awarded: 2004

URL: <http://tinyurl.com/yvkbfo>

Abstract

Context-aware systems exploit the use of situational information, or context, to provide relevant information and services to users. A great challenge remains in defining an architecture that supports context-aware systems. Critical research issues include modeling and reasoning (how to represent contextual information for machine processing and reasoning), knowledge sharing (how to enable agents to acquire consistent knowledge from unreliable sensors and agents), and user privacy protection (how to give users control of their private information that the system acquires).

To address these issues, I developed a new agent architecture called the Context Broker Architecture (CoBrA). It uses the Web Ontology Language (OWL) to define ontologies for context representation and modeling, defines rule-based logical inference for context reasoning and knowledge maintenance, and provides a policy language for users to control the sharing of their private information. Central to CoBrA is a server agent called context broker. Its role is to maintain a consistent model of context that can be shared by all computing entities in the space and to enforce the user-defined policies for privacy protection.

The major research contributions of this work include a broker-centric architecture for supporting context-

aware systems, a standard pervasive computing ontology, a reasoning approach that integrates assumption-based reasoning and argumentation for resolving inconsistent contextual knowledge, and a privacy protection mechanism that exploits information granularity adjustment.

To demonstrate the feasibility of CoBrA, I prototyped a context broker in the FIPA platform using the JADE API library. I showed its use in supporting EasyMeeting, a smart meeting room system that provides context-aware services for assisting speakers and audiences. Other contributions include the CoBrA Demo Toolkit (an open source software package for demonstrating various aspects of CoBrA) and the CoBrA Text Messaging Commands (a text messaging interface for mobile users to interact with a context broker via SMS messages).

The lessons learned from this research are as the follows. (i) CoBrA's broker-centric design can help to reduce the time and effort to rapidly prototype context-aware applications. (ii) Ontologies expressed using the OWL language can provide a uniformed solution for context representation and reasoning, knowledge sharing, and meta-language definitions. (iii) Rule-based logical inference can help to develop flexible context-aware systems by separating high-level context reasoning from low-level system behaviors.

A Principled Methodology for the Design of Autonomous Trading Agents with Combinatorial Preferences in the Presence of Tradeoffs

Candidate: Ioannis A. Vetsikas

Institution: Computer Science Department, Cornell University, Ithaca, NY, USA

Supervisor: Bart Selman

Year awarded: 2005

Abstract

Online auctions have become a popular method for business transactions. The variety of different auction rules, the restrictions in supply or demand, and the agents' combinatorial preferences for the different commodities, have led to the creation of a very complex multi-agent 'environment' and a number of

strategic tradeoffs. Designing an agent that deals efficiently with these tradeoffs has been a multi-pronged effort. Using game-theoretic approaches, some equilibria have been computed for relatively simple auctions. However, since these equilibria have limited practical application, due to the significant

number of varying auctions that take place simultaneously, empirical approaches and experimental evaluations of various strategies have also been used. Furthermore, progress has been made into designing better agent architectures. This dissertation presents results in all of these directions (theoretical and empirical). We present a methodology for designing trading agents, and deciding their bidding strategy, when they participate in a large number of simultaneous auctions with a variety of rules. We use a modular, adaptive, scalable and robust agent architecture, combining principled methods and empirical knowledge. We decompose the problem faced by the agent into several components, and restrict communication between them. This allows us to analyze each component individually. The ‘optimizer’ coordinates the effort of the entire system and determines the set of commodities that maximizes the utility of the agent. We provide a principled way of generating strategies for each bidding module. Then we use rigorous experimentation to explore the

strategy space and determine the best combination of strategies. This allows the agent to perform efficiently against any opponent agents. We also expand this methodology to include design decisions based on the equilibria computed for particular auctions.

Furthermore, we present several novel Bayes–Nash equilibria for m^{th} price multi-unit auctions with multiple possible closing times, one of which is chosen randomly, and therefore multiple rounds of bidding can occur.

We applied this methodology when creating WhiteBear, the agent that won most of the Trading Agent Competitions (TACs) held between 2001 and 2005 and had the best performance overall. We also present the ‘complete’ set of experiments for determining an overall best strategy in TAC, which guided the design of WhiteBear. We show that exploring the strategy space allows the creation of more efficient and more flexible agents than any other approach, e.g. using learning.

Exploring Complex Dynamics in Multi Agent–Based Intelligent Systems: Theoretical and Experimental Approaches Using the Multi Agent–Based Behavioral Economic Landscape (MABEL) Model

Candidate: Konstantinos T. Alexandridis

Institution: Department of Forestry and Natural Resources, Purdue University, IN, USA

Supervisor: Bryan C. Pijanowski

Year awarded: 2006

URL: <http://docs.lib.purdue.edu/dissertations/AAI3232142/>

Abstract

This dissertation adopts a holistic and detailed approach to modeling spatially explicit agent-based artificial intelligent systems, using the Multi Agent–based Behavioral Economic Landscape (MABEL) model. The research questions that addresses stem from the need to understand and analyze the real-world patterns and dynamics of land use change from a coupled human–environmental systems perspective. Describes the systemic, mathematical, statistical, socio-economic and spatial dynamics of the MABEL modeling framework, and provides a wide array of cross-disciplinary modeling applications within the research, decision-making and policy domains. Establishes the symbolic properties of the MABEL model as a Markov decision process, analyzes the decision-theoretic utility and optimization attributes of agents toward comprising statistically and spatially optimal policies and actions, and explores the probabilistic character of the agents’ decision-making and inference mechanisms via the use of Bayesian belief and decision networks. Develops and describes

a Monte Carlo methodology for experimental replications of agent’s decisions regarding complex spatial parcel acquisition and learning. Recognizes the gap on spatially-explicit accuracy assessment techniques for complex spatial models, and proposes an ensemble of statistical tools designed to address this problem. Advanced information assessment techniques such as the Receiver–Operator Characteristic curve, the impurity entropy and Gini functions, and the Bayesian classification functions are proposed. The theoretical foundation for modular Bayesian inference in spatially-explicit multi-agent artificial intelligent systems, and the ensembles of cognitive and scenario assessment modular tools build for the MABEL model are provided. Emphasizes the modularity and robustness as valuable qualitative modeling attributes, and examines the role of robust intelligent modeling as a tool for improving policy-decisions related to land use change. Finally, the major contributions to the science are presented along with valuable directions for future research.

Dynamically Self-reconfigurable Systems for Machine Intelligence

Candidate: Haibo He

Institution: School of Electrical Engineering and Computer Science, Ohio University, Athens, OH, USA

Supervisor: Janusz A. Starzyk

Year awarded: 2006

URL: <http://www.ece.stevens-tech.edu/~hhe/>

Abstract

This dissertation is focused on the development of system level architectures and models of dynamically self-reconfigurable systems for machine intelligence. This research is significant for building brain-like intelligent systems. Although the development of deep submicron very large scale integration (VLSI) system, nanotechnology and bioinformatics facilitate building such intelligent systems, yet it is very challenging to study how these kinds of complex, reconfigurable systems can self-develop their connectivity structures, accumulate knowledge, make associations and predictions, dynamically interact with environment, and self-control to accomplish desired tasks.

A new framework of 'learning-memory-prediction' for machine intelligence is proposed in this research, and it serves as the foundation for building intelligent systems through learning in dynamic value systems, memorizing in self-organizing networks, and predicting in hierarchical structures. These systems are characterized by on-line data-driven learning, distributed structure of processing components with local and sparse interconnections, dynamic reconfigurability, self-organization, and active interaction with environment.

Learning is the fundamental element for biologically intelligent systems. The proposed online value system

is able to learn and dynamically estimate the value of any multi-dimensional data set, and such value system can be used in reinforcement learning. Feedback mechanism is introduced in the self-organizing learning system to allow the machine to be able to memorize information in its distributed processing elements and make associations. After the information is learned and stored in the associative memory, a biologically inspired anticipation-based temporal sequence learning architecture is proposed. All systems proposed in this research are hardware oriented. A novel computing paradigm that can achieve low power consumption for designing large scale, high density intelligent systems is proposed, and a brief description of the system level hardware architecture for prototyping and testing of the proposed systems is also presented.

Intelligent systems have wide applications from military security systems to civilian daily life. In this research, different application problems, including pattern recognition, classification, image recovery, and sequence learning, are presented to show the capability of the proposed systems in learning, memory, and prediction.

Abstraction, Aggregation and Recursion for Generating Accurate and Simple Classifiers

Candidate: Dae-Ki Kang

Institution: Department of Computer Science Iowa State University, Ames, IA, USA

Supervisor: Vasant Honavar

Year awarded: 2006

URL: <http://www.cs.iastate.edu/~dkkang/>

Abstract

An important goal of inductive learning is to generate accurate and compact classifiers from data. In a typical inductive learning scenario, instances in a data set are simply represented as ordered tuples of attribute values. In my research, I explore three methodologies to improve the accuracy and compactness of the classifiers: abstraction, aggregation, and recursion. First, *abstraction* is aimed at the design and analysis of algorithms that generate and deal with taxonomies for the construction of compact and robust classifiers. In many applications of the data-driven knowledge discovery process, taxonomies have been shown to be useful in constructing compact, robust, and comprehensible classifiers. However, in many application domains, human-designed taxonomies are unavailable. I introduce algorithms for automated construction of taxonomies inductively from both structured (such as UCI Repository) and unstructured (such as text and biological sequences) data. I introduce AVT-Learner, an algorithm for automated construction of attribute value taxonomies (AVT) from data, and Word Taxonomy Learner (WTL), an algorithm for automated construction of word taxonomy from text and sequence data. I describe experiments on the UCI data sets and compare the performance of AVT-NBL (an AVT-guided Naive Bayes Learner) with that of the standard Naive Bayes Learner (NBL). The results show that the AVTs generated by AVT-Learner are competitive with human-generated AVTs (in cases where such AVTs are available). AVT-NBL using AVTs generated by AVT-Learner achieves classification accuracies that are comparable to or higher than those obtained by NBL;

and the resulting classifiers are significantly more compact than those generated by NBL. Similarly, the experimental results of WTL and WTNBL on protein localization sequences and Reuters newswire text categorization data sets show that the proposed algorithms can generate Naive Bayes classifiers that are more compact and often more accurate than those produced by standard Naive Bayes learner for the Multinomial Model.

Second, I apply *aggregation* to construct features as a multiset of values for the intrusion detection task. For this task, I propose a bag of system calls representation for system call traces and describe misuse and anomaly detection results on the University of New Mexico (UNM) and MIT Lincoln Lab (MIT LL) system call sequences with the proposed representation. With the feature representation as input, I compare the performance of several machine learning techniques for misuse detection and show experimental results on anomaly detection. The results show that standard machine learning and clustering techniques using the simple bag of system calls representation based on the system call traces generated by the operating system's kernel is effective and often performs better than approaches that use foreign contiguous sequences in detecting intrusive behaviors of compromised processes. Finally, I construct a set of classifiers by *recursive application* of the Naive Bayes learning algorithms. Naive Bayes (NB) classifier relies on the assumption that the instances in each class can be described by a *single* generative model. This assumption can be restrictive in many real world classification tasks. I describe recursive Naive Bayes learner (RNBL),

which relaxes this assumption by constructing a tree of NB classifiers for sequence classification, where each individual NB classifier in the tree is based on an event model (one model for each class at each node in the tree). In my experiments on protein sequences, Reuters newswire documents and UC-Irvine benchmark data sets, I observe that RNBL substantially outperforms

NB classifier. Furthermore, my experiments on the protein sequences and the text documents show that RNBL outperforms C4.5 decision tree learner (using tests on sequence composition statistics as the splitting criterion) and yields accuracies that are comparable to those of support vector machines (SVM) using similar information.

On Mental Probability Logic

Candidate: Niki Pfeifer

Institution: Department of Psychology, University of Salzburg, Austria

Supervisor: Gernot D. Kleiter

Year awarded: 2006

URL: <http://www.users.sbg.ac.at/~pfeifern/>

Abstract

Mental probability logic is a psychological competence theory about how humans interpret and reason about common-sense conditionals. Probability logic is proposed as an appropriate standard of reference for evaluating the rationality of human inferences. Common-sense conditionals are interpreted as high conditional probabilities, $P(B|A) > 0.5$. Probability logic accounts of

non-monotonic reasoning and inference rules like the MODUS PONENS are explored. Categorical syllogisms with comparative and quantitative quantifiers are investigated. A series of eight experiments on human probabilistic reasoning in the framework of the basic nonmonotonic SYSTEM P corroborate the psychological plausibility of the proposed approach.

Evaluating the Efficacy of Mathematically Modelling the Behaviour of Engineered Systems

Candidate: T. Eshan Rajabally

Institution: Faculty of Science, Agriculture and Engineering, University of New-castle upon Tyne, Newcastle, UK

Supervisor: Pratyush Sen

Year awarded: 2006

URL: <http://www-staff.lboro.ac.uk/~ouer/index.html>

Abstract

Mathematical modelling plays a huge role in analyzing the behaviour of engineered systems. Despite the benefits of modelling, inappropriate or indefensible applications can have unwanted and even disastrous consequences. The fundamental reason for such indefensible applications is the pervasive role of human judgement throughout any modelling endeavour and the subsequent scope for uncertainty. This role is compounded by two key difficulties of modern day modelling. First, there is a lack of objective means for considering the numerous criteria that a model must fulfil. Secondly, modelling is prone to uncertainties and errors and, due to our incomplete understanding of the world, the absolute fidelity of a model cannot be demonstrated. As a result, evaluating the efficacy of mathematical modelling is essentially an exercise in human reasoning under uncertainty and current means of efficacy evaluation fail to capture this reasoning.

A novel approach to model efficacy evaluation based upon Bayesian Belief Networks is proposed in this thesis that captures the key role of human reasoning

under uncertainty. Means to reduce the task of quantifying a network are suggested, coupled with means to reduce the influence of human subjectivity in the resulting approach to model efficacy evaluation. For the purpose of the thesis, an assessment of the approach is performed in two ways. First, the steps of the approach are illustrated on a hypothetical network of interoperating models. A critique of the approach including advantages and disadvantages is thus provided without divulging commercially sensitive details of proprietary models. Second, an analysis of two defence-based industrial case studies of the approach are provided that concern the modelling of military air-craft survivability and naval operational capability. The assessment of the approach demonstrates how use of the Bayesian Belief Network approach satisfies requirements for a more holistic evaluation of model efficacy than can currently be performed. The key benefit of Bayesian updating is the understanding it generates of the efficacy of networks of models such as suites, hierarchies and synthetic environments.

Learning Interpretable Models

Candidate: Stefan Rüping

Institution: Computer Science Department, University of Dortmund, Germany

Supervisor: Katharina Morik

Year awarded: 2006

URL: <http://hdl.handle.net/2003/23008>

Abstract

Interpretability is an important, yet often neglected criterion when applying machine learning algorithms to real-world tasks. An understandable model enables the user to gain more knowledge from his data and to participate in the knowledge discovery process in a more detailed way. Hence, learning interpretable models is a challenging task, whose complexity comes from the problems that interpretability is a fuzzy, subjective concept and human mental capabilities are in some ways astonishingly limited. At the same time, interpretability is a critical problem, because it is crucial for problems that cannot be solved purely automatically. The work presented in this thesis is structured along the three dimensions of understandability, accuracy, and efficiency. It contains

contributions on the levels of the optimization of the interpretability of a learner with and without knowledge of its internals (white box and black box approach), the description of a models errors by local patterns and the improvement of global models with local models. Starting from an analysis of the requirements for and measures of interpretability in the context of knowledge discovery, diverse possible approaches of generating understandable models are investigated, with a particular focus on interpretable Support Vector Machines and local effects in the data. Problems of existing techniques and ad-hoc approaches to understandability optimization are analyzed and improved algorithms are developed.

The Structure and Behaviour of the Continuous Double Auction

Candidate: Perukrishnen Vytelingum

Institution: School of Electronics and Computer Science, University of Southampton, Southampton, UK

Supervisor: Nicholas R. Jennings

Year awarded: 2006

URL: <http://www.ecs.soton.ac.uk/~pv03r>

Abstract

The last decade has seen a shift in emphasis from centralized to decentralized systems to meet the demanding coordination requirements of today's complex computer systems. In such systems, the aim is to achieve effective decentralized control through autonomous software agents that perform local decision-making based on incomplete and imperfect information. Specifically, when the various agents interact, the system behaves as a computational ecology with no single agent coordinating their actions. In this thesis, we focus on one specific type of computational ecology, the Continuous Double Auction (CDA), and investigate market-oriented approaches to decentralized control. In particular, the CDA is a fixed-duration auction mechanism where multiple buyers and sellers compete to buy and sell goods, respectively, in the market, and where transactions can occur at any time whenever an offer to buy and an offer to sell match. Now, in such a market mechanism, the decentralized control is achieved through the decentralized allocation of resources, which, in turn, is an emergent behaviour of buyers and sellers trading in the market. The CDA was chosen, among the plenitude of auction formats available, because it allows efficient resource allocation without the need of a centralized auctioneer. Against this background, we look at both the structure and the behaviour of the CDA in our attempt to build an efficient and robust mechanism for decentralized control. We seek to do this for both stable environments, in which the market demand and supply do not change and dynamic ones in which there are sporadic changes (known as market shocks). While the structure of the CDA defines the agents' interactions in the market, the behaviour of the CDA is determined by what emerges when the buyers and sellers compete to maximize their individual profits. In more detail, on the structural aspect, we first look at how the market protocol of the CDA can be modified to meet desirable properties for the system (such as

high market efficiency, fairness of profit distribution among agents and market stability). Second, we use this modified protocol to efficiently solve a complex decentralized task allocation problem with limited-capacity suppliers that have start-up production costs and consumers with inelastic demand. Furthermore, we demonstrate that the structure of this CDA variant is very efficient by evaluating the mechanism with very simple agent behaviours. In so doing, we emphasise the effect of the structure, rather than the behaviour, on efficiency.

In the behavioural aspect, we first developed a multi-layered framework for designing strategies that autonomous agents can use for trading in various types of market mechanisms. We then use this framework to design a novel Adaptive-Aggressiveness (AA) strategy for the CDA. Specifically, our bidding strategy has both a short and a long-term learning mechanism to adapt its behaviour to changing market conditions and it is designed to be robust in both static and dynamic environments. Furthermore, we also developed a novel framework that uses a two-population evolutionary game theoretic approach to analyze the strategic interactions of buyers and sellers in the CDA. Finally, we develop effective methodologies for evaluating strategies for the CDA in both homogeneous and heterogeneous populations, within static and dynamic environments. We then evaluate the AA bidding strategy against the state-of-the-art using these methodologies. By so doing, we show that, within homogeneous populations, the AA strategy outperformed the benchmarks, in terms of market efficiency. Within heterogeneous populations, based on our evolutionary game-theoretic framework, we identify that the AA strategy will eventually be more likely to be adopted by buyers and sellers in the market (for being more efficient).

Insertion of Virtual Players in Role-Playing Games to use in Group Decision Support Systems: A Case Study in Natural Resources Management

Candidate: Diana Francisca Adamatti

Institution: Departamento de Engenharia de Computação e Sistemas Digitais, Escola Politécnica da Universidade de São Paulo, Brasil

Supervisor: Jaime Simo Sichman

Year awarded: 2007

URL: www.lti.pcs.usp.br/~adamatti

Abstract

The thesis proposes the definition of an architecture to insert virtual players in a particular subclass of Group Decision Support Systems (GDSS), that uses Multi-Agent-Based Simulation (MABS) and Role-Playing Games (RPGs) techniques in an integrated way. These techniques can bring interesting results, since it is possible to join the dynamic capacity of MABS with the discussion and learning capacity of RPGs. The defined virtual players must make decisions and communicate with each other and with the real players during the negotiation process. In this work, the main

aspects discussed are the following: (i) can virtual players have non-trivial behavior in the face of real players? (ii) is the negotiation process between all players (virtual or real) still happening when virtual players are inserted? (iii) do electronic games make more difficult the interaction between players? In order to discuss these aspects, two instances of GDSSs were developed and tested in the natural resources management domain. This domain was chosen because its negotiation process is both very important and complex.

Scalable Planning Under Uncertainty

Candidate: Daniel Bryce

Institution: Department of Computer Science and Engineering, Arizona State University, Tempe, AZ, USA

Supervisor(s): Subbarao Kambhampati

Year awarded: 2007

URL: <http://verde.eas.asu.edu/dissertation.html>

Abstract

Autonomous agents that act in the real-world can often improve their success by capturing the uncertainty that arises because of their imperfect knowledge and potentially faulty actions. By making plans robust to uncertainty, agents can be prepared to counteract plan failure or act upon information that becomes available during plan execution. Such robust plans are valuable, but are often difficult to construct when uncertainty is high and goal achievement requires many actions. Unfortunately, current approaches to planning under uncertainty are not scalable: they focus on small problems where the challenge is finding an optimal solution. I study (non-deterministic and probabilistic) conformant and conditional planning in large problems where just

finding a feasible solution is the challenge. I develop scalable heuristic search techniques that are inspired both by optimal approaches to planning under uncertainty (based on Markov decision processes) and scalable approaches for classical planning (based on planning graph reachability heuristics). Specifically, I develop measures for the cost of completing partial plans, heuristics that estimate the measures, efficient data-structures and inference techniques to compute the heuristics, and a multi-objective heuristic search algorithm that uses the heuristics. Through extensive empirical evaluation, I show the resulting set of techniques significantly improves the scalability of planning under uncertainty.

Optimizing and Implementing Repair Programs for Consistent Query Answering in Databases

Candidate: Monica Caniupan

Institution: School of Computer Science, Carleton University, Ottawa, ON, Canada

Supervisor: Leopoldo Bertossi

Year awarded: 2007

URL: <http://www.face.ubiobio.cl/~mcaniupa/index.html>

Abstract

Databases may not always satisfy their integrity constraints (ICs) and a number of different reasons can be held accountable for this. However, in most cases an important part of the data is still consistent with the ICs, and can still be retrieved through queries posed to the database. Consistent query answers are characterized as ordinary answers obtained from every minimally repaired and consistent version of the database. Database repairs

wrt a wide class of ICs can be specified as stable models of disjunctive logic programs. Thus, Consistent Query Answering (CQA) for first-order queries is translated into cautious reasoning under the stable models semantics. The use of logic programs does not exceed the intrinsic complexity of CQA. However, using them in a straightforward manner is usually inefficient. The goal of this thesis is to develop optimized techniques to

evaluate queries over inconsistent databases by using logic programs. More specifically, we optimize the structure of programs, model computation, and evaluation of queries from them. We develop a system which implements optimized logic programs and efficient methods to compute consistent answers to first-order queries.

Moreover, we propose the use of the well-founded semantics (WFS) as an alternative way to obtain

consistent answers. We show that for a certain class of queries and ICs, the well-founded interpretation of a program retrieves the same consistent answers as the stable models semantics. The WFS has lower data complexity than the stable models semantics.

We also extend the use of logic programs for retrieving consistent answers to aggregate queries, and we develop a repair semantics for Multidimensional Databases.

Reasoning with Dynamic Networks in Practice

Candidate: Theodoros Charitos

Institution: Department of Information and Computing Sciences, University of Utrecht, the Netherlands

Supervisor: Linda C. van der Gaag

Year awarded: 2007

URL: <http://www.cs.uu.nl/staff/theodore.html>

Abstract

The modelling and analysis of time-evolving phenomena constitutes a significant topic in science and engineering. Dynamic Bayesian networks are powerful and flexible graphical models for representing and computing probability distributing over variables that relate to stochastic processes. For a set of variables capturing phenomena that evolve over time, such models specify a set of conditional independence assumptions that allows the joint distribution to be represented in a factored way. Efficient inference with such models, however, remains a crucial issue for their successful application in practice.

The main contribution of this thesis lies in methods for easing inference with dynamic networks by exploitation of either the nature of the data or the parameters of the model. To provide for a realistic setting, the thesis first considers the extension of a static model for the management of patients at intensive care units suspected of suffering from ventilator-associated pneumonia, into a dynamic network. Using this dynamic network and the real-life data accompanying it as the main vehicle upon

which all methods are constructed, this thesis demonstrates the effect of the model's parameters on its output probability distribution or on a decision for antibiotic treatment based upon this distribution. The precise form of mathematical functions describing these effects is established, while approximate methods for efficient computing of these functions are also proposed. Next, the thesis presents flexible inference algorithms that are tailored to the application at hand and exploit either consecutive similar values from diagnostic tests or symptoms of several patients that are sequentially observed, or the specifications of the probabilities in the transition matrices and in the sensitivity and specificity rates of several diagnostic tests of the model. Finally, the thesis presents an algorithm for efficient inference in the case where the observed data arrive at arbitrary points in time instead of with pre-defined transition intervals. As a result, the distribution of the hidden variables is approximated at such time by interpolating between the boundaries of the pre-defined intervals while using the algorithm for further computations.

Bayesian Networks for Clinical Decision Support

Candidate: Marcel van Gerven

Institution: Institute for Computing and Information Sciences, Radboud University, Nijmegen, the Netherlands

Supervisors: Peter Lucas and Theo van der Weide

Year awarded: 2007

URL: <http://osiris.cs.ru.nl/~marcelge/>

Abstract

The dissertation handles about the use of (dynamic) Bayesian networks for the support of physicians with respect to decision making in clinical oncology. The research focuses on the manual construction of such networks and the automated learning of such networks using a limited amount of data. The dissertation subsequently deals with methods for the manual construction of medical Bayesian networks, a

qualitative analysis of causal independence models, the approximate solution of partially observable Markov decision processes, the development of a large dynamic Bayesian network for treatment of carcinoid tumours, and the analysis of three novel classification algorithms, namely, the maximum mutual information algorithm, the decomposed tensor classifier, and the noisy-threshold classifier.

Increasing Accessibility in Agent-Oriented Methodologies

Candidate: Jorge Gonzalez-Palacios

Institution: School of Electronics and Computer Science, University of Southampton, Southampton, UK

Supervisor: Michael Luck

Year awarded: 2007

Abstract

The computing world is rapidly changing from a one in which a centralized approach is taken to one in which a highly distributed approach is taken, thus requiring software systems that operate in open, dynamic and heterogeneous environments. This has significantly increased the complexity of software systems, and has required the development of new paradigms for software development, such as the multi-agent approach to system development. However, even though there is evidence of the suitability of the multi-agent approach to cope with the complexity of current systems, its use is not widespread in other areas of computing science, nor in industrial and commercial environments. This can be explained, particularly for agent-oriented methodologies, by the absence of key software engineering best practices. In particular, we have identified three groups of drawbacks that limit the use of agent-oriented methodologies: incomplete coverage of the development cycle, a lack of tools for supporting the development process, and a high degree of dependence on specific toolkits, methods or platforms. Although these issues negatively affect the applicability of the multi-agent approach in general, it is arguably for open systems that their effect is particularly noticeable.

In this thesis, therefore, we aim to address the issues involved in taking existing agent-oriented methodologies to a point where they can be effectively applied to the development of open systems. In order to do so, we consider the combination of organizational design and agent design, as well as the methodological process itself. Specifically, we address organizational design by constructing a software engineering technique (software patterns) for the representation and incorporation of standard organizations into the organizational design of a multi-agent system. The agent design aspect is addressed by constructing an agent design phase which uses standard agent architectures through a pattern catalogue. Based on this, we develop a methodological process that combines the organizational and agent designs, and that also considers the use of iterations for making the development of a system more agile. This methodological process is exemplified and assessed by means of a case study. Finally, we address the problem of monitoring the correct behaviour of agents in an open system, by constructing a model for the specification of open multi-agent systems.

Uncertainty and Indistinguishability. Application to Modelling with Words

Candidate: Enric Hernndez Jimnez

Institution: Llenguatges i Sistemes Informtics Department, Universitat Politcnica de Catalunya (UPC), Barcelona, Catalunya, Spain

Supervisor: Jordi Recasens Ferrs

Year awarded: 2007

URL: <http://www.tdx.cesca.es/TDX-0419107-100926/index.html>

Abstract

The concept of equality is a fundamental notion in any theory since it is essential to the ability of discerning the objects to whom it concerns, ability which in turn is a requirement for any classification mechanism that might be defined.

When all the properties involved are entirely precise, what we obtain is the classical equality, where two individuals are considered equal if and only if they share the same set of properties. What happens, however, when imprecision arises as in the case of properties which are fulfilled only up to a degree? Then, because certain individuals will be more similar than others, the need for a gradual notion of equality arises.

These considerations show that certain contexts that are pervaded with uncertainty require a more flexible concept of equality that goes beyond the rigidity of the classic concept of equality. T-indistinguishability operators seem to be good candidates for this more flexible and general version of the concept of equality that we are searching for.

On the other hand, Dempster-Shafer Theory of Evidence, as a framework for representing and managing general evidences, implicitly conveys the notion of indistinguishability between the elements of

the domain of discourse based on their relative compatibility with the evidence at hand. In chapter two, we are concerned with providing definitions for the T-indistinguishability operator associated to a given body of evidence.

In chapter three, after providing a comprehensive summary of the state-of-the-art on measures of uncertainty, we tackle the problem of computing entropy when an indistinguishability relation has been defined over the elements of the domain. Entropy should then be measured not according to the occurrence of different events, but according to the variability perceived by an observer equipped with indistinguishability abilities as defined by the indistinguishability relation considered. This idea naturally leads to the introduction of the concept of observational entropy.

Real data is often pervaded with uncertainty so that devising techniques intended to induce knowledge in the presence of uncertainty seems entirely advisable. The paradigm of computing with words follows this line in order to provide a computation formalism based on linguistic labels in contrast to traditional numerical-based methods.

The use of linguistic labels enriches the understandability of the representation language, although it also requires adapting the classical inductive learning procedures to cope with such labels. In chapter four, a novel approach to building decision trees is introduced, addressing the case when uncertainty arises as a consequence of considering a more realistic setting in which decision maker's discernment abilities are taken into account when computing node's impurity measures. This novel

paradigm results in what have been called 'observational decision trees' since the main idea stems from the notion of observational entropy in order to incorporate indistinguishability concerns.

In addition, we present an algorithm intended to induce linguistic rules from data by properly managing the uncertainty present either in the set of describing labels or in the data itself. A formal comparison with standard algorithms is also provided.

Managing Discourse and Uncertainty for Decision-making in Civil and Infrastructure

Engineering Systems

Candidate: Emad Marashi

Institution: Civil Engineering, University of Bristol, UK

Supervisor(s): John P. Davis

Year awarded: 2007

Abstract

This research attempts to develop an integrated methodology and tool that brings about a step forward in managing discourse and uncertainties for decision-making. The decisions and actions of engineers have a profound impact on the world and society at large. As such, it is important that engineers communicate and justify their decision-making process and reasoning path to the stakeholders involved. These tasks can be supported and further improved through a systemic development of appropriate methodologies and tools that would help in modelling complex systems, structuring the problem, facilitating dialogue, and managing knowledge and uncertainties. Process modelling has been used to provide the means for gaining a holistic view of the system. The success of each process depends significantly on its ability to address issues that arise, which could be subject to discourse and negotiation. This debate and discussion has been represented in the Issue-Based Information System framework for argumentation. The imperfect,

incomplete and vague nature of information highlights the need for soft approaches to reasoning with uncertainty. More generic methods for combination and propagation of imprecise and uncertain evidence are thus developed in this research. The Evidential Discourse for Engineering (EDEN) is introduced as an integrated framework for process and argumentation modelling as a result, which is implemented in the PeriMeta-2 software tool. The application of the EDEN methodology to the case of climate change impacts on the UK electricity supply industry has facilitated the management of complex and contentious arguments into a relatively simple, but evidence-based, graphical model. This helps the strategic decision makers to gain a broader view on how the performance of their systems is affected in different demand and supply scenarios, and how interventions and various adaptation options can impact their processes and the industry as a whole.

Probabilistic Models with Unknown Objects

Candidate: Brian Christopher Milch

Institution: Computer Science Division, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, CA, USA

Supervisor(s): Stuart J. Russell

Year awarded: 2007

URL: <http://people.csail.mit.edu/milch/blog>

Abstract

Humans and other intelligent agents must make inferences about the real-world objects that underlie their observations: for instance, the objects visible in an image, or the people mentioned in a set of text documents. The agent may not know in advance how many objects exist, how they are related to each other, or which observations correspond to which underlying objects. Existing declarative representations for probabilistic models do not capture the structure of such scenarios.

This thesis introduces Bayesian logic (BLOG), a first-order probabilistic modeling language that specifies probability distributions over possible worlds with varying sets of objects. A BLOG model contains statements that define conditional

probability distributions for a certain set of random variables; the model also specifies certain context-specific independence properties. We provide criteria under which such a model is guaranteed to fully define a probability distribution. These criteria go beyond existing results in that they can be satisfied even when the Bayesian network defined by the model is cyclic, or contains nodes with infinitely many ancestors.

We describe several approximate inference algorithms that exploit the context-specific dependence structure revealed by a BLOG model. First, we present rejection sampling and likelihood weighting algorithms that are guaranteed to converge to the correct probability for any query on a structurally well-defined BLOG model.

Because these algorithms instantiate only those variables that are context-specifically relevant, they can generate samples in finite time even when the model defines infinitely many variables. We then define a general framework for inference on BLOG models using Markov chain Monte Carlo (MCMC) algorithms. This framework allows a programmer to plug in a domain-specific proposal distribution, which helps the Markov chain move to high-probability

worlds. Furthermore, the chain can operate on partial world descriptions that specify values only for context-specifically relevant variables. We give conditions under which MCMC over such partial world descriptions is guaranteed to converge to correct probabilities. We also show that this framework performs efficiently on a real-world task: reconstructing the set of distinct publications referred to by a set of bibliographic citations.

New Representations and Approximations for Sequential Decision Making under Uncertainty

Candidate: Tao Wang

Institution: Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada

Supervisors: Dale Schuurmans and Michael Bowling

Year awarded: 2007

URL: <http://www.cs.ualberta.ca/~trysi>

Abstract

This dissertation research addresses the challenge of scaling up algorithms for sequential decision making under uncertainty. In my dissertation, I developed new approximation strategies for planning and learning in the presence of uncertainty while maintaining useful theoretical properties that allow larger problems to be tackled than is practical with exact methods. In particular, my research tackles three outstanding issues in sequential decision making in uncertain environments: performing stable generalization during off-policy updates, balancing exploration with exploitation, and handling partial observability of the environment. The first key contribution of my thesis is the development of novel dual representations and algorithms for planning and learning in stochastic environments. This dual view I have developed offers a coherent and comprehensive approach to optimal sequential decision making problems, provides a viable alternative to standard value function based techniques for solving dynamic programming and reinforcement learning problems. In particular, I have shown that dual dynamic programming algorithms can avoid the divergence problems associated with the standard primal approach, even in the presence of approximation and off-policy updates. Another key contribution of my thesis is the development of a practical action selection strategy

that addresses the well known exploration versus exploitation tradeoff in reinforcement learning. The idea is to exploit information in a Bayesian posterior to make intelligent actions by growing an adaptive, sparse lookahead tree. This technique evaluates actions while taking into account any effects they might have on future knowledge, as well as future reward, and outperforms current selection strategies. Finally, my thesis also develops a new approach to approximate planning in partially observable Markov decision processes. Here the challenge is to overcome the exponential space required by standard value iteration. For this problem, I introduced a new, quadratic upper bound approximation that can be optimized by semidefinite programming. This approach achieves competitive approximation quality while maintaining a compact representation; requiring computational time and space that is only linear in the number of decisions. Overall, my dissertation research developed new tools for computing optimal sequential decision strategies in stochastic environments, and has contributed significant progress on three key challenges in reinforcement learning.