

Abstracts of Recent PhDs

Regression in the Presence of Incomplete Information and Sensing Actions, and its Application to Conditional Planning

Candidate: Tuan Chi Le

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

Supervisor: Chitta Baral

Year awarded: 2004

Abstract

In this dissertation, a state-based regression function for planning domains where an agent does not have complete information and may have sensing actions is presented. Both binary and multi-valued domains are considered, and the 0-approximation [SB01] is employed to define regression with respect to that semantics. In binary domains, the use of 0-approximation means using three-valued (true, false, and unknown) states. In multi-valued domains, each fluent in a state is assigned an unknown value or a value in a finite set of the fluent's prescribed values. Although planning using this approach is incomplete with regard to the full semantics, it is adopted to have a lower complexity. The soundness and completeness of the regression formulation with regard to

the definition of progression are presented. More specifically, the dissertation shows that a plan obtained through regression for a planning problem is indeed a progression solution of that planning problem, and that for each plan found through progression, using regression one obtains that plan or an equivalent one. A contingent planner that utilizes the regression function is then developed and the soundness and completeness of the planning algorithm are proved. Heuristic measures are also employed to improve the planning performance. Experimental results with respect to several well-known planning problems in the literature and self-created domains are presented.

Computational Intelligence Methods on Biomedical Signal Analysis and Data Mining in Medical Records

Candidate: Liviu Mihai Vladutu

Institution: Department of Medical Physics, University of Patras, Greece

Supervisors: Anastasios Bezerianos, Michael Vrahatis and Nicholas Fakotakis

Year awarded: 2004

URL: http://heart.med.upatras.gr/bio/people/Liviu_phd_Thesis.pdf

Abstract

This thesis is centred around the development and application of computationally effective solutions based on artificial neural networks (ANN) for biomedical signal analysis and data mining in medical records. The goal of this work in the field of Biomedical Engineering is to provide the clinician with the best possible information needed to make an accurate diagnosis (in our case of myocardial ischaemia) and to propose advanced mathematical models for recovering the complex dependencies between the variables of a physical process from a set of perturbed observations.

We designed a model for pattern classification, by constructing several local models, for neighbourhoods of the state space. For this task, we use the novel k-windows clustering algorithm, to automatically detect neighbourhoods in the state space. This algorithm, slightly modified (unsupervised k-windows algorithm), has the ability to endogenously determine the number of clusters present in the dataset during the clustering

process. We used this method together with the other two mentioned below (NetSOM and sNet-SOM) for the problem of ischaemia detection. We continue by establishing the network self-organizing map (NetSOM) model, which attempts to generalize the regularization and ordering potential of the basic SOM from the space of vectors to the space of approximating functions. It becomes a device for the ordering of local experts (independent neural networks) over its lattice of neurons and for their selection and co-ordination.

Finally, an alternative to NetSOM is proposed, which uses unsupervised ordering based on self-organizing maps (SOM) for the 'simple' regions and for the 'difficult' ones a two-stage learning process. In comparison with the previous model (NetSOM), we replaced a fixed size of the SOM with a dynamically expanded map and, second, the supervised learning was based this time on RBF networks and SVM. This tool was used in two fields, namely ischaemia detection and data mining.

Intelligent Tutoring for Medical Problem-Based Learning

Candidate: Siriwan Suebnukarn

Institution: Faculty of Dentistry, Thammasat University, Thailand

Supervisor: Peter Haddawy

Year awarded: 2005

URL: http://www.dentistry.tu.ac.th/Dentistry/Page_05/Text_01/Siriwan.html

Abstract

This dissertation describes intelligent tutoring in a collaborative medical tutor for problem-based learning (PBL). The main contribution of this work is the development of individual and collaborative student modeling techniques and algorithms for generating tutoring hints in PBL group problem solving, as well as the implementation of these techniques in a collaborative intelligent tutoring system, COMET. The system combines concepts from computer-supported collaborative learning with those from intelligent tutoring systems. COMET is designed to provide an experience that emulates that of live human-tutored medical PBL sessions as much as possible while at the same time permitting the students to participate collaboratively from disparate locations. COMET incorporates a multi-modal interface that integrates text and graphics so as to provide a rich communication channel between the students and the system, as well as among students in the group. Generating appropriate tutorial actions in

COMET requires a model of the students' clinical reasoning for the problem domain. In addition, since problem solving in group PBL is a collaborative process, modeling individuals and the group is necessary if we wish to develop an intelligent tutoring system that can do things like focus the group discussion, promote collaboration, or suggest peer helpers. Medical PBL is particularly challenging due to the complexity of the knowledge involved, the lack of standard, commonly accepted student clinical-reasoning techniques, and the lack of standards for tutoring. This means that we must first attempt to identify prototypical patterns of clinical reasoning and then formalize them to create the clinical reasoning model. The system uses Bayesian networks to model individual student clinical reasoning, as well as that of the group. This dissertation describes general domain-independent individual and collaborative clinical reasoning models and tutoring algorithms for intelligent medical tutoring in group settings.

Enhancing Spatial Association Rule Mining in Geographic Databases

Candidate: Vania Bogorny

Institution: Instituto de Informática, Universidade Federal Do Rio Grande Do Sul (UFRGS), Brazil

Supervisors: Luis Otavio Alvares and Paulo Martins Engel

Year awarded: 2006

URL: www.inf.ufrgs.br/vbogorny

Abstract

The association rule mining technique emerged with the objective to find novel, useful, and previously unknown patterns from non-spatial data. The main drawback of this technique is the generation of large amounts of frequent patterns and association rules. In geographic databases, the problem of mining spatial association rules increases significantly. Besides the large amount of generated patterns and rules, many are redundant or well-known geographic domain associations explicitly represented in geographic database schemas and geo-ontologies. The result is that the same associations explicitly represented in such knowledge repositories are extracted by spatial association rule mining algorithms and presented to the user. The problem of mining spatial association rules from geographic databases requires at least three main steps: compute spatial relationships, generate frequent patterns, and extract association rules. The first step is the most effort demanding and time-consuming task in the rule mining process, but has received little attention in the literature. The second and third steps have been considered the main problem in non-spatial association rule mining and have been addressed as two different problems:

frequent pattern mining and association rule mining. Well-known geographic dependences that generate well-known patterns may appear in the three main steps of the spatial association rule mining process. Aiming to eliminate both redundant and well-known dependences in spatial association rule mining, this thesis presents a general framework with three main novel methods for mining frequent geographic patterns using background knowledge. The first method reduces the input space, and all well-known dependences that can be eliminated without losing information are removed in data preprocessing. The second method is an early pruning strategy, which eliminates combinations of pairs of geographic objects with dependences, during the frequent set generation. The third method is a new approach for eliminating both redundant and well-known patterns, in order to generate maximal generalized frequent geographic patterns without dependences. Several experiments with real geographic databases have shown that the proposed framework efficiently reduces the total number of geographic patterns, and as a consequence, generates more interesting association rules.

Large-scale Information Acquisition for Data and Information Fusion

Candidate: Ronnie Johansson

Institution: The Royal Institute of Technology, Sweden

Supervisors: Henrik I. Christensen, Stefan Arnborg and Per Svensson

Year awarded: 2006

URL: http://www.csc.kth.se/cgi-bin/ronniej/gr_counter/thesis.cgi

Abstract

The purpose of information acquisition for data and information fusion is to provide relevant and timely information. The acquired information is integrated (or fused) to estimate the state of some environment. The success of information acquisition can be measured in the quality of the environment state estimates generated by the data and information fusion process.

In this thesis, we introduce and set out to characterize the concept of large-scale information acquisition. Our interest in this subject is justified both by the identified lack of research on a holistic view on data and information fusion and by the proliferation of networked sensors that promises to enable handy access to a multitude of information sources. We identify a number of properties that could be considered in the context of large-scale information acquisition. The sensors used could be large in number, heterogeneous, complex, and distributed. Also, algorithms for large-scale information acquisition may have to deal with decentralized control and multiple and varying objectives. In the literature, a process that realizes information acquisition is frequently denoted sensor management. We, however, introduce the term perception management instead, which encourages an agent perspective on information acquisition. Apart from explicitly inviting the wealth of agent theory research into the data and information fusion research, it also highlights that the resource usage of perception management is constrained by the overall control of a system that uses data and information fusion. To address the challenges posed by

the concept of large-scale information acquisition, we present a framework that highlights some of its pertinent aspects. We have implemented some important parts of the framework. What becomes evident in our study is the innate complexity of information acquisition for data and information fusion, which suggests approximative solutions.

We, furthermore, study one of the possibly most important properties of large-scale information acquisition, decentralized control, in more detail. We propose a recurrent negotiation protocol for (decentralized) multi-agent coordination. Our approach to the negotiations is from an axiomatic bargaining theory perspective: an economics discipline. We identify shortcomings of the most commonly applied bargaining solution and demonstrate in simulations a problem instance where it is inferior to an alternative solution. However, we cannot conclude that one of the solutions dominates the other in general. They are both preferable in different situations. We have also implemented the recurrent negotiation protocol on a group of mobile robots.

We note some subtle difficulties with transferring bargaining solutions from economics to our computational problem. For instance, the characterizing axioms of solutions in the bargaining theory are useful to qualitatively compare different solutions, but care has to be taken when translating the solution to algorithms in computer science, as some properties might be undesirable, unimportant, or at risk of being lost in the translation.

Visual Novelty Detection for Autonomous Inspection Robots

Candidate: Hugo Vieira Neto

Institution: Department of Computer Science, University of Essex, UK

Supervisor: Ulrich Nehmzow

Year awarded: 2006

URL: <http://pessoal.cefetpr.br/hugo/pub.html>

Abstract

Mobile robot applications that involve automated exploration and inspection of environments are often dependant on novelty detection, the ability to differentiate between common and uncommon perceptions. Because novelty can be anything that deviates from the normal context, we argue that in order to implement a novelty filter it is necessary to exploit the robot's sensory data from the ground up, building models of normality rather than abnormality.

In this work, we use unrestricted colour visual data as perceptual input to on-line incremental learning algorithms. Unlike other sensor modalities, vision can provide a variety of useful information about the environment through massive amounts of data, which often need to be reduced for real-time operation. Here, we

use mechanisms of visual attention to select candidate image regions to be encoded and fed to higher levels of processing, enabling the localization of novel features within the input image frame.

An extensive series of experiments using visual input, obtained by a real mobile robot interacting with laboratory and medium-scale real-world environments, are used to discuss different visual novelty filter configurations. We compare performance and functionality of novelty detection mechanisms based on the Grow-When-Required neural network and incremental principal component analysis. Results are assessed using both qualitative and quantitative methods, demonstrating the advantages and disadvantages of each investigated approach.

Sketch Interpretation Using Multiscale Stochastic Models of Temporal Patterns

Candidate: Tevfik Metin Sezgin

Institution: Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA

Supervisor: Randall Davis

Year awarded: 2006

URL: <http://rationale.csail.mit.edu/publications/Sezgin2006Sketch.pdf>

Abstract

Sketching is a natural mode of interaction used in a variety of settings. For example, people sketch during early design and brainstorming sessions to guide the thought process; when we communicate certain ideas, we use sketching as an additional modality to convey ideas that cannot be put in words. The emergence of hardware such as PDAs and Tablet PCs has enabled capturing freehand sketches, enabling the routine use of sketching as an additional human-computer interaction modality.

But despite the availability of pen-based information capture hardware, relatively little effort has been put into developing software capable of understanding and reasoning about sketches. To date, most approaches to sketch recognition have treated sketches as images (i.e. static finished products) and have applied vision algorithms for recognition. However, unlike images, sketches are produced incrementally and interactively, one stroke at a time and their processing should take advantage of this.

We explore the ways of doing sketch recognition by extracting as much information as possible from temporal

patterns that appear during sketching. We present a sketch recognition framework based on hierarchical statistical models of temporal patterns. We show that in certain domains, stroke orderings used in the course of drawing individual objects contain temporal patterns that can aid recognition. We build on this work to show how sketch recognition systems can use knowledge of both common stroke orderings and common object orderings. We describe a statistical framework based on dynamic Bayesian networks that can learn temporal models of object-level and stroke-level patterns for recognition. Our framework supports multi-object strokes and multi-stroke objects, and allows interspersed drawing of objects—relaxing the assumption that objects are drawn one at a time. Our system also supports real-valued feature representations using a numerically stable recognition algorithm. We present recognition results for hand-drawn electronic circuit diagrams. The results show that modeling temporal patterns at multiple scales provides a significant increase in correct recognition rates, with no added computational penalties.

A Generative Framework for Argumentation-Based Enquiry Dialogs

Candidate: Elizabeth Black

Institution: Advanced Computation Laboratory, Cancer Research, UK; Computer Science Department, University College London, UK

Supervisors: Anthony Hunter and John Fox

Year awarded: 2007

URL: <http://www.acl.icnet.uk/lab/phdLiz.html>

Abstract

This PhD focusses on argumentation-based communication between agents. I take, as a starting point, the argumentation system proposed by Garcia and Simari [2004, *Defeasible logic programming an argumentative approach. Theory and Practice of Logic Programming* 4(1–2), 95–138], which allows a single agent to reason about its beliefs. I define a novel dialogue system that allows two agents to use Garcia and Simari's system to carry out inter-agent argumentation. I define two specific protocols for two different types of enquiry dialogue that I define: argument enquiry and warrant enquiry. Argument enquiry dialogues are often embedded within warrant enquiry dialogues. Other existing enquiry dialogue systems only model dialogs, meaning that they describe what a legal enquiry dialogue is, but they do not provide the means to actually generate such a dialogue. Such systems provide a protocol, which dictates what the possible legal next moves are at each point in a dialogue

but not which of these moves to make. I present a system that not only includes two dialogue game style protocols, one for the argument enquiry dialogue and one for the warrant enquiry dialogue, but also includes an intelligent strategy, for an agent to use with these protocols, that selects exactly one of the legal moves to make. As my system is generative, it allows me to investigate the precise behaviour of the dialogues it produces. I propose a benchmark against which I compare my dialogues, and use this to define soundness and completeness properties for argument enquiry and warrant enquiry dialogues. I show that these properties hold for all dialogues produced by my system. Finally, I go on to define another intelligent strategy for use with warrant enquiry dialogues. I show that this also leads to sound and complete dialogues but, in many situations, reduces the redundancy seen in the dialectical tree produced during the dialogue.

Handling Inconsistency in Databases and Data Integration Systems

Candidate: Loreto Bravo

Institution: School of Computer Science, Carleton University, Canada

Supervisor: Leopoldo Bertossi

Year awarded: 2007

URL: <http://homepages.inf.ed.ac.uk/lbravo/>

Abstract

For several reasons, a database may not satisfy certain integrity constraints (ICs), for example, when it is the result of integrating several independent data sources. However, most likely, information in it is still consistent

with the ICs, and could be retrieved when queries are answered. Consistent answers with respect to a set of ICs have been characterized as answers that can be obtained from every possible minimal repair of the database. The

goal of this research is to develop methods to retrieve consistent answers for a wide and practical class of constraints and queries from relational databases and from data integration systems. We will put special interest on databases with null values. We will give a semantics of satisfaction of constraints in the presence of null that generalizes the one used in commercial DBMS. Since there are interesting connections between the area of

consistently querying virtual data integration systems and other areas, like querying incomplete databases, merging inconsistent theories, semantic reconciliation of data, schema mapping, data exchange, and query answering in peer data management systems, the results of this research could also be applied to them. In our research, we explore in more depth the connection with virtual data integration systems and peer data management systems.

Speciated Evolutionary Modularization and Integration for Real-World Applications

Candidate: Kyung-Joong Kim

Institution: Department of Computer Science, Yonsei University, Republic of Korea

Supervisor: Sung-Bae Cho

Year awarded: 2007

URL: <http://candy.yonsei.ac.kr/uribyl>

Abstract

Modular approaches are popular because they can provide the bottom-up style construction of complex systems. However, the modularization is usually done manually by human experts and this becomes a bottleneck for more popular use of the approaches. In biology, the modularization process at the genome level proceeds through an evolutionary procedure. The function and characteristics of modules are changed by genetic operations, and they are combined to form a large-scale complex system. The separation of modules in large systems is a self-organizing procedure based on the interaction of building blocks. Similarly, the modularization in artificial complex systems can be automated using evolutionary computation.

Integration of multiple modules has high complexity because of the myriad of possible combinations. To solve this problem, the evolutionary approach, which is highly effective in global search, has been used. The automatic integration and modularization have been performed by evolutionary computation. However, they can easily get stuck in local optima; also, the integration of similar modules may not improve the performance. In this thesis, a speciation algorithm is used in the evolutionary

modularization and integration to overcome these problems. The speciation method maintains diverse solutions, which are regarded as modules. Also, it accelerates the convergence of searching integrator. Finally, we extend previous work by demonstrating its application to real-world applications.

Usually, real-world applications require preprocessing, incorporation of domain knowledge, and uncertainty handling. In this thesis, we deal with issues in the framework of speciated evolutionary modularization and integration. For preprocessing, a number of different feature extraction methods are automatically attached to classification modules using a speciated evolutionary algorithm. Domain knowledge is incorporated into the evaluation process of speciated modularization, to reduce the search space. Finally, Bayesian networks are used to deal with uncertainty in real-world applications.

The proposed framework is evaluated on applications in bioinformatics, games, and knowledge discovery. Each issue of real-world application is tested separately, using one of the three problems. Experimental results show that the proposed framework can improve the performance of speciated modularization and integration.

Incorporating Prior Domain Knowledge into Inductive Machine Learning—Its Implementation in Contemporary Capital Markets

Candidate: Ting Yu

Institution: Faculty of Information Technology, University of Technology, Sydney, Australia

Supervisors: John Debenham, Tony Jan, Simeon Simoff and Donald Stokes

Year awarded: 2007

URL: <http://www.physics.usyd.edu.au/tingyu/thesis.pdf>

Abstract

An ideal inductive machine-learning algorithm produces a model best approximating an underlying target function by using reasonable computational cost. This requires the resultant model to be consistent with the training data, and generalize well over the unseen data. Regular inductive machine-learning algorithms rely heavily on numerical data as well as on general-purpose inductive bias. However, certain environments contain rich domain knowledge prior to the learning task, but it is not easy for regular inductive learning algorithms to utilize prior domain knowledge. This thesis discusses and analyses

various methods of incorporating prior domain knowledge into inductive machine learning through three key issues: consistency, generalization, and convergence. Additionally, three new methods are proposed and tested over data sets collected from capital markets. These methods utilize financial knowledge collected from various sources, such as experts and research papers, to facilitate the learning process of kernel methods (emerging inductive learning algorithms). The test results are encouraging and demonstrate that prior domain knowledge is valuable to inductive learning machines.