

Abstracts of Recent PhDs

Qualified Predictions for Large Data Sets

Candidate: Harris Papadopoulos

Institution: Computer Science Department of Royal Holloway, University of London, England, UK

Supervisors: Alex Gammerman and Volodya Vovk

Year awarded: 2004

doi:10.1017/S0269888909990221

Abstract

The Inductive Confidence Machine (ICM) provides an alternative method to that of the Transductive Confidence Machine (TCM) for complementing the bare predictions produced by traditional machine-learning algorithms with measures of confidence. These measures give an indication of how ‘good’ each prediction is, which is highly desirable in risk-sensitive applications. The motivation behind the introduction of the ICM was to produce algorithms that overcome the computational inefficiency problems suffered by TCMs.

In this thesis, we study the ICM method, describing how it works and how it can be applied to different traditional machine-learning algorithms. More

specifically we detail how we implemented the Nearest Neighbours and Neural Networks ICMs for pattern recognition, and the Ridge Regression and Nearest Neighbours Regression ICMs. The results obtained by our methods demonstrate that the accuracy of ICMs in terms of error percentage is comparable to both traditional methods and TCMs and that the confidence measures they produce are useful in practice. In addition, our time efficiency comparisons exhibit their huge advantage in this sector over TCMs. These properties make ICMs the most suitable choice for obtaining qualified predictions when dealing with large data sets.

Monitoring Large-Scale Multi-Agent Systems using Overhearing

Candidate: Gery Gutnik

Institution: Computer Science Department, Bar-Ilan University, Israel

Supervisor: Gal A. Kaminka

Year awarded: 2006

URL: <http://www.cs.biu.ac.il/~gutnikg>

doi:10.1017/S0269888909990233

Abstract

Overhearing is fast gaining attention as a generic method for monitoring open, distributed multi-agent systems. In such settings, agents’ internal structure is not generally known to a monitoring agent, but overhearing does not require such knowledge. Instead, the monitoring agent uses the overheard routine communications as a basis for inference about the other agents.

Previous work on overhearing investigated an extensive set of techniques and implementations of overhearing. However, focusing mainly on its potential applications, those investigations often rely on assumptions related to the fundamentals of overhearing. In contrast, we dedicate our research to a comprehensive study of the fundamental building blocks that allow overhearing in the first place tackling those problematic assumptions. In particular, our study focuses on overhearing in large-scale multi-agent systems and addresses the specific challenges and limitations that characterize such settings.

The first overhearing building block, addressed by our research, is the representation of multi-agent conversations. Here, building on the insights gained from analyzing the strengths and weaknesses of the rather

radical Petri net approaches introduced by previous work, we propose a novel representation technique especially suitable for overhearing. Furthermore, we show this representation to be more scalable than previous representations, and thus more appropriate for monitoring conversations in large-scale settings.

Next, we addressed the building block of conversation recognition—the process of identifying the actual conversation based on a sequence of overheard messages. Although conversation recognition is a key step in overhearing prior to any possible inference, it is often discarded by previous investigations. Our work addresses the challenges related to conversation recognition by first introducing a formal model of overhearing. Then, based on this model, we provide a skeleton algorithm for conversation recognition, and provide instantiations of it for lossless and lossy settings. Since in large-scale settings overhearing agent has to process large quantities of intercepted messages, we also analyze the efficiency of those algorithms in terms of their run-time complexity.

The final building block addressed is selective overhearing, that is, overhearing under the restriction of

selectivity. The restriction of selectivity is mainly compelled by the specific characteristics of large-scale multi-agent systems. In such settings, it is reasonable to assume that the overhearing resources will be essentially limited, thus allowing the overhearing agent to overhear only a subset of inter-agent communications. Most

previous investigations on overhearing ignore the limitation of selectivity, assuming that all relevant inter-agent communications can be overheard. In contrast, our work provides an empirical study of selective overhearing committed by both centralized and distributed teams of collaborative overhearing agents.

Cognitively Inspired Decision Making for Software Agents: Integrated Mechanisms for Action Selection, Expectation, Automatization and Non-Routine Problem Solving

Candidate: Aregahegn Seifu Negatu

Institution: Computer Science Department, The University of Memphis, TN, USA

Supervisor: Stan Franklin

Year awarded: 2006

URL: <http://ccrg.cs.memphis.edu/>

doi:10.1017/S0269888909990245

Abstract

Despite impressive advances in the past decades, autonomous agents living in dynamic and unpredictable environments are typically equipped with simple decision-making mechanisms in their sense-decide-act routines. These agents deal mostly with one goal at a time. This research aspires to model, design and/or implement a sophisticated decision-making mechanism that selects the agent's next action with different levels of awareness: automatized skills, consciously mediated routine solutions and consciously deliberated non-routine solutions. Such a decision-making mechanism is presented in a 'conscious' software agent framework called IDA that implements Baars' Global Workspace Theory of consciousness. IDA integrates many computational and conceptual mechanisms, among which this research deals with its action selection, expectation, automatization and non-routine problem-solving modules.

The overarching continual task of an agent's intelligence is for the service of choosing, at each moment in time, the appropriate action in response to exogenous and endogenous stimuli. IDA's action selection mechanism (ASM) can interleave and prioritize actions of different and competing

goal hierarchies. The ASM system is implemented as a domain-independent and reusable framework for behavior networks and is tested as a controller to a khepera robot operating in a real-world domain.

We humans have the amazing ability to learn a procedural task (e.g. walking) so well that we do not need to think about the task consciously in order to accomplish it. This ability is what we call automatization. Once a task has been automatized, there is no need for attention to be paid to its execution unless the expected result does not occur. At failure of expectation, deautomatization process temporarily disables the automatization effects and 'conscious' control plays a role to deal with the failure situation. We implement the automatization and deautomatization cognitive functions as a self-organizing system in the IDA framework.

Non-routine problem solving is the ability to devise unexpected, and often clever, solutions to problems that have never been encountered before. We will present a detailed design and specification of a non-routine problem-solving mechanism as a special goal context hierarchy that guides a deliberative solution search process, which we will discuss in IDA's cognitive cycle.

Autonomy vs. Conformity: An Institutional Perspective on Norms and Protocols

Candidate: Huib Aldewereld

Institution: Intelligent Systems Group, Institute of Computer and Information Sciences, Universiteit Utrecht, The Netherlands

Supervisors: John-Jules Meyer and Frank Dignum

Year awarded: 2007

URL: <http://igitur-archive.library.uu.nl/dissertations/2007-0604-200758/>

doi:10.1017/S0269888909990257

Abstract

The research presented in this thesis is part of the ToKeN project ANITA (Administrative Normative Information Transaction Agents). In this domain, the main challenges concern both the shortage of information (not being able to find legally relevant data that should be available) as well as the abundance of information (e.g., violating privacy rights). Although the agents decide autonomously whether to share information based on (local) norms, a global frame was needed for the enforcement of (global) norms given by laws and regulations of the domain. In most software and agent methodologies, such regulations are seen only as extra requirements in the analysis phase, and are thus hard-coded into the software or agents themselves. If, however the regulations change, all design steps have to be checked and all code verified to ensure

compliance to the new regulations. The alternative is to have an explicit representation of the norms, for example, by the introduction of an electronic institution, but this approach requires some form of enforcement to ensure the compliance instead. The introduction of an electronic institution in highly regulated domains such as the ANITA scenario requires us to solve issues related to the abstractness of human regulations, the lack of operational information and the implementation of norm enforcement from an institutional perspective. In this thesis, we solve these problems by the introduction of a framework for making the connections between the norms and the practice explicit. To ensure that none of the norms derived from law are violated, an enforcement mechanism is presented based on monitoring and punishing rather than