

Abstracts of Recent PhDs

What Should We Do? Computational Representation of Persuasive Argument in Practical Reasoning

Candidate: Katie Marie Atkinson

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

Supervisors: Peter McBurney and Trevor Bench-Capon

Year awarded: 2005

URL: www.csc.liv.ac.uk/katie

Abstract

The design and development of autonomous software agents requires a multitude of elements to be considered and accounted for. For a software agent to be considered ‘intelligent’, it must be able to perform effective reasoning about its beliefs and the environment in which it is situated, and also act in this environment. It must therefore reason about both what is the case and what should be done: the latter is known as practical reasoning. In addition, it must also be able to interact and reason with other such agents in its environment, as it may rely on them for information and help to enable it to accomplish its tasks. This thesis is concerned with one particular aspect of such agency: modelling the process of argument in practical reasoning to equip autonomous agents with the ability to determine the best action to take in a given situation. The background setting for this work deals with the topic of practical reasoning and attempts to address some issues regarding its treatment in philosophy, as well as the problems inherent in the computational modelling of such reasoning. The main output of the study is a theory of persuasion in practical reasoning which makes use of techniques from the field of argumentation theory to enable autonomous software agents to construct and reason about arguments in

support of and against proposals for action. The theory is embodied in a model describing how agents based on the belief–desire–intention (BDI) architecture can put forward a proposal for action and how this proposal can be systematically attacked in a variety of ways. This enables agents to consider all available options and come to a conclusion about the best action to take in the given context. The underlying theory extends a well-established account from the field of philosophy based on the use of argument schemes and critical questions. The account given is then formalized to enable its representation in agent systems. The underlying theory has formed the basis for a number of applications: an implementation of a dialogue game protocol to provide a proof of concept, an implementation to provide computer-mediated support for human decision-making in a particular context and, finally, a formalism to enable autonomous agents to reason about decisions regarding actions. The account for use in BDI agents is applied to three example domains—law, medicine and politics—to show how BDI agents can reason and argue about matters of practical action, in accordance with the theory.

Programming Uncertain Agents

Candidate: Nivea de Carvalho Ferreira

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

Supervisors: Michael Fisher and Wiebe van der Hoek

Year awarded: 2006

URL: www.csc.liv.ac.uk/niveacf

Abstract

Agent-based systems have been widely investigated in recent years. As they provide unambiguous semantics, formal languages appear to be an appropriate formalization for an agent’s specification. Logical formalization of an agent’s behaviour is desirable not only to provide a clear semantics of agent-based systems but also to provide the foundation for sophisticated reasoning techniques. Moreover, the use of a formal language brings the possibility of

implementing the agent’s behaviour by direct execution of logical statements, serving as a link between semantics and practical implementation of agents. The possible-worlds semantics offered by modal logic has proved to be a successful framework to model mental attitudes of agents such as beliefs, desires and intentions. In addition, when referring to the ‘real world’, we usually deal with inaccurate information which is not definitely believed. In such scenarios, making use of formalisms that allow

the appropriate representation of uncertainty (to some level and to some extent) becomes essential. We intend to show, throughout this thesis, the development (and application) of a logic-based programming language for describing and animating an agent's dynamic behaviour. In essence, ProTem is a powerful, although simple, logical language obtained by the combination of a linear temporal logic-based framework, MetateM,

with a probabilistic logic of belief, PF KD45. Among other accomplishments, we propose new mechanisms for handling uncertainty in executable temporal doxastic logic specifications. Potential application areas include planning and scheduling, communication protocols and games with mixed strategies.

Fault-Tolerant Knowledge-Level Inter-Agent Communication in Open Multi-Agent Systems

Candidate: Nicola Dragoni

Institution: Department of Computer Science, University of Bologna, Italy

Supervisor: Mauro Gaspari

Year awarded: 2006

URL: www.cs.unibo.it/dragon/

Abstract

Agent communication languages (ACLs) have been developed to provide a way for agents to communicate with each other, supporting cooperation in multi-agent systems. The importance of ACLs is especially highlighted in open systems (such as the Internet) where heterogeneous agents must be able to communicate and exchange knowledge with one another despite differences in hardware platforms, operating systems, architectures and programming languages. In the past few years, many ACLs have been proposed for multi-agent systems, such as knowledge query manipulation language and Foundation for Intelligent Physical Agents' ACL. The goal of these languages is to support high-level, human-like communication among agents, exploiting knowledge-level features rather than symbol-level ones. Despite these efforts, one important issue in the research on ACLs is still open and concerns how these languages should deal with failures of agents in open

multi-agent systems. The main consequence arising from this lack is that it is not possible to develop knowledge-level agents that are able to cooperate in an open multi-agent system prone to agent failures. The fault-tolerant agent communication language (FT-ACL) presented in this thesis addresses this issue. FT-ACL provides high-level fault-tolerant communication primitives and supports for an anonymous interaction protocol which satisfies a set of well-defined knowledge-level programming requirements. We present a formal specification of FT-ACL and the underlying agent architecture. To illustrate the language features, we show how FT-ACL can be effectively used to write high-level executable specifications of fault-tolerant protocols. Moreover, we show how FT-ACL can be used to support Web agent interaction as well as Web Service invocation on the SemanticWeb.

Argumentation-Based Negotiation in a Social Context

Candidate: Nishan C. Karunatilake

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

Supervisor: Nicholas R. Jennings

Year awarded: 2006

URL: www.ecs.soton.ac.uk/people/nnc

Abstract

Argumentation-based negotiation (ABN) is gaining importance as a fundamental method of interaction in multi-agent systems. In essence, ABN enhances the ways agents can interact within a negotiation encounter. In particular, it allows agents to justify their demands, criticize each other's proposals and add comments to their statements during a negotiation encounter. Furthermore, ABN gives agents the ability to exchange explicit arguments, such as promises, threats, appeals and other forms of persuasive locutions, to influence one another during a negotiation dialogue. Such enhancements lead to richer forms of negotiation than have hitherto been possible in game-theoretic or heuristic-based models. Therefore, many argue that endowing agents with the ability to argue during their negotiation interactions not only facilitates more realistic rational dialogues but also allows an effective means of resolving different forms of conflicts endemic

to multi-agent societies. Even though ABN is argued to be an effective means of resolving conflicts, its operation within multi-agent systems incurs certain computational overheads. In particular, it takes time for an agent to argue and convince an opponent to change its demands and yield to a less favourable agreement within an ABN encounter. It also takes computational effort for both parties to the conflict to carry out the reasoning required to generate, select and evaluate an appropriate and convincing set of arguments required for such an encounter. On the other hand, within a multi-agent society, not all conflicts need to be resolved. In some instances, conflicts can be avoided by other, non-arguing means. For instance, in certain situations, agents may be able to avoid conflicts by finding an alternative resource to achieve their actions instead of arguing over a conflicting one. They may also be able to re-plan to

achieve the same objective through a different means and, thereby, remove the conflict without argument. In the presence of such overheads and given the alternatives available, this thesis argues that computationally bounded entities such as agents need to consider two critical questions before they use ABN to manage their conflicts. The first is when to argue; that is, under what conditions would ABN, as opposed to other, non-arguing methods, present a better option for agents to overcome conflicts? The second is how to argue, that is, a computationally tractable method and a set of strategies to formulate such sophisticated ABN dialogues successfully. To this end, this thesis advances a detailed theoretical and empirical study to address both these research questions. In more detail, first we formulate a novel ABN framework that allows agents to argue, negotiate and, thereby, resolve conflicts in structured multi-agent systems. The framework is unique in the way that it explicitly captures social influences endemic to such agent societies and, in turn, allows agents to use them constructively in their ABN dialogues. Having formulated the framework, we then map it onto the computational context of a multi-agent

task allocation scenario. In so doing, we bridge the gap between theory and practice and provide a test-bed to evaluate how our ABN model can be used to manage and resolve conflicts in multi-agent societies. Our experimental analysis on when to argue shows a clear inverse correlation between the benefit of arguing and the resources available within the context. It also shows that arguing selectively is both a more efficient and a more effective strategy than doing so in an exhaustive manner. Furthermore, we show that when agents operate under imperfect knowledge conditions, an arguing approach allows them to perform more effectively than a non-arguing one. On the issue of how to argue, we show that arguing earlier in an ABN interaction presents a more efficient method than arguing later in the interaction. Moreover, during an ABN interaction, allowing agents to negotiate their social influences presents both an effective and an efficient method which will enhance their performance within a society.

On the Mutability of Protocols

Candidate: Jarred McGinnis

Institution: School of Informatics, University of Edinburgh, UK

Supervisor: David Robertson

Year awarded: 2006

URL: www.era.lib.ed.ac.uk/handle/1842/1403

Abstract

This thesis addresses the dichotomy between reliable communication and facilitation of the autonomy of agents to create more flexible and emergent interactions. By the introduction of adaptations to a distributed protocol language, agents benefit from the ability to communicate interaction protocols to elucidate the social norms. Yet, this approach also provides the functionality for the agent to unilaterally introduce new paths for the conversation to explore unforeseen opportunities and options. The foundation of this work is lightweight coordination calculus (LCC). LCC is a distributed protocol language and framework in which agents coordinate their own interactions by their message passing activities. To ensure that adaptations to the protocols are made in a reasonable way, we examine the use of two models of communication to guide any transformations to the protocols. We describe the use of Foundation for Intelligent Physical Agents' agent communication language and ultimately its unsuitability for this approach as well as the more

fecund task of implementing dialogue games, and models of argumentation, as dynamic protocols. The existing attempts to develop a model that can encompass the gulf between reliability and autonomy in communication have had varying degrees of success. The purpose and the result of the research described in this thesis is to develop an alloy of the various models, by the introduction of dynamic and distributed protocols, to develop a framework stronger than its constituents. Though this is successful, the derivations of the protocols can be difficult to reconstruct. To this end, this thesis also describes a method of protocol synthesis inspired by models of human communication that can express the dialogues created by the previous approaches and that also have a fully accountable path of construction. This thesis explores a unique and novel approach to agent communication and also tests it through a practical implementation.

Agent-Based Trust and Reputation in the Context of Inaccurate Information Sources

Candidate: W. T. Luke Teacy

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

Supervisors: Nicholas R. Jennings and Michael Luck

Year awarded: 2006

URL: www.ecs.soton.ac.uk/wtlt/

Abstract

Trust is a prevalent concept in human society that, in essence, concerns our reliance on the actions of other entities within our environment. For example, we may rely on our car starting to get to work on time, and on

our fellow drivers so that we can get there safely. For similar reasons, trust is becoming increasingly important in computing, as systems such as the Grid require integration of computing resources across

organizational boundaries. In this context, the reliability of resources in one organization cannot be assumed from the point of view of another, as certain resources may fail more often than others. For this reason, we argue that software systems must be able to assess the reliability of different resources so that they can choose which of them to rely on. With this in mind, our goal is to develop mechanisms, or models, to aid decision-making by an autonomous agent (the truster) when the consequences of its decisions depend on the actions of other agents (the trustees). To achieve this, we have developed a probabilistic framework for assessing trust based on a trustee's past behaviour, which we have instantiated through the creation of two novel trust models, trust and reputation model for agent-based virtual organizations (TRAVOS) and TRAVOS-C. These facilitate decision-making in two different contexts with regard to trustee behaviour. First, using TRAVOS, a truster can make decisions in contexts where a trustee can act only in one of two ways: either it can cooperate, acting to the truster's

advantage, or it can defect, thereby acting against the truster's interests. Second, using TRAVOS-C, a truster can make decisions about trustees that can act in a continuous range of ways, for example, taking into account the delivery time of a service. These models share an ability to account for observations of a trustee's behaviour made either directly by the truster or by a third party (reputation source). In the latter case, both models can cope with third-party information that is unreliable, either because the sender is lying or because it has a different worldview. In addition, TRAVOS-C can assess a trustee for which there is little or no direct or reported experience, using information on other agents that share characteristics with the trustee. This is achieved using a probabilistic mechanism which automatically accounts for the amount of correlation observed between agents' behaviour in a truster's environment.
