An Integrated Framework for Adaptive Reasoning About Conversation Patterns*

Michael Rovatsos
mrovatso@inf.ed.ac.uk
School of Informatics
University of Edinburgh
Edinburgh EH8 9LE, UK

Felix Fischer
fischerf@cs.tum.edu
Department of Informatics
Technical University of Munich
85748 Garching, Germany

Gerhard Weiss
weissg@cs.tum.edu
Department of Informatics
Technical University of Munich
85748 Garching, Germany

ABSTRACT

We present an integrated approach for reasoning about and learning conversation patterns in multiagent communication. The approach is based on the assumption that information about the communication language and protocols available in a multiagent system is provided in the form of dialogue sequence patterns, possibly tagged with logical conditions and instance information. We describe an integrated social reasoning architecture m^{2}InFFrA that is capable of (i) processing such patterns, (ii) making communication decisions in a boundedly rational way, and (iii) learning patterns and their strategic application from observation.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems, Languages and Structures*

General Terms

Languages, Theory

Keywords

Agent Communication, Evolutionary Semantics, Social Reasoning, Interaction Frames

1. INTRODUCTION

Compared to the long-established areas of interaction protocol and agent communication language (ACL) research [1], the development of agent architectures suitable for dealing with given communication mechanisms in practical terms has received fairly little attention. As yet, there exists no uniform framework for defining the interface between the inter-agent communication layer and intra-agent reasoning, i.e. how specifications of interaction protocols and communication semantics influence agent rationality or, in turn, are influenced themselves by agents' rational decision-making.

*A full version of this paper is available at `http://homepages.inf.ed.ac.uk/mrovatso/papers/rovatsosetal-aamas2005-full.pdf`

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands

In this paper, we attempt to tackle this problem from a very pragmatic perspective. We make very weak assumptions regarding the method used to define the available means of communication in a multiagent system (MAS), namely that it provides (i) a description of the surface structure of communication processes (in the simplest case, traces of possible message and action sequences in agent conversations) that is tied to (ii) some form of logical constraints (in a tractable logical language, if they are to be used in agent reasoning). In the following, we refer to such pairs of surface structure and logical constraints as *conversation patterns*.

2. INTERACTION FRAMES

The greatest common denominator of the multitude of methods for specifying ACL semantics and interaction protocols is that they describe the *surface structure* of dialogues (i.e. a set of admissible message sequences) and logical *constraints* for the applicability of these message sequences (which may include statements about environmental conditions, mental states of the participating agents, the state of commitment stores, etc.). In the most simplistic case, these structure/constraint pairs may be represented as combinations of a conversation trace and a set of logical conditions. The m^{2}InFFrA architecture [6] we describe here uses *interaction frames* to represent such patterns and augments them with frequency counters that allow for the definition of a probabilistic semantics. Consider the following example of such a frame:

$$F = \Big\langle \big\langle \xrightarrow{5} \texttt{request}(A,B,X) \xrightarrow{3} \texttt{do}(B,X) \big\rangle, \\ \big\langle \{can(B,X)\}, \{can(B,\texttt{pay}(S)\} \big\rangle \\ \big\langle \xrightarrow{2} \langle [A/\texttt{a}],[B/\texttt{b}],[X/\texttt{pay}(\$100)] \rangle, \\ \xrightarrow{1} \langle [A/\texttt{b}],[B/\texttt{a}],[X/\texttt{pay}(S)] \rangle \big\rangle \Big\rangle$$

This frame reflects the following interaction experience: A has asked B five times to perform (physical) action X, B actually did so in three of these instances. In two of the successful instances, it was a who asked and b who headed the request, and the action was to pay \$100. In both cases, $can(\texttt{b},\texttt{pay}(\$100))$ held true. In the third case, roles were swapped between a and b and the amount S remains unspecified (which does not mean that it did not have a concrete value, but that this information was abstracted away in the frame).

An important feature of m^{2}InFFrA frames in contrast to general conversation patterns is that they allow for storing *empirical* information about past conversation instances that followed a certain pattern and also to distinguish between different sets of conditions that held during these *enactments* of a frame.

3. REASONING WITH FRAMES

The ability of frames to capture instance information enables agents to reason about communication semantics in an adaptive fashion. In accordance with the *empirical semantics* view [8] that considers the meaning of communication as a function of its consequences as experienced through the eyes of a subjective observer, agents can adapt existing frame conceptions from new observations and project past regularities into the future. As we will see, this can improve their strategic communication abilities decisively, particularly in *open* systems where agents may or may not obey predefined conversation patterns.

3.1 Frame Semantics

According to the *probabilistic interpretation* of frames in the m^{2}InFFrA model, the semantics of (a set of) frames is defined as follows: Given an *encounter prefix* w, i.e. a (possibly empty) sequence of messages already uttered in the current encounter and a *knowledge base* KB of beliefs the reasoning agent currently holds, we can compute the set of possible *continuations* w' (i.e. message sequences that will conclude the current encounter) by (i) filtering out all those frames whose trajectories do not prefix-match w, (ii) considering the postfixes of w in the remaining frames under the remaining possible substitutions (given that w has already committed certain variables to specific values), and (iii) applying those substitutions whose corresponding condition sets are satisfied under KB.

3.2 Decision Making with Frames

If agents were equipped with a *utility estimate* $u(w, KB) \in \mathbb{R}$ that allows them to assess the usefulness of a particular sequence w of messages (and actions) in belief state KB, they could in principle sum up continuation probabilities over all frames to derive utility-maximising decisions. However, this not only contradicts our goal of breaking down the whole network of communicative expectations held by an agent into manageable "chunks" (i.e. frames), but also the way in which conversation patterns for MASs are usually defined (i.e. in terms of different protocols for different purposes, not all of which need constantly be reasoned over while engaging in a particular kind of interaction).

Instead we will assume (in a boundedly rational fashion) that an agent only activates a single frame at a time within which it then searches for an optimal action while engaging in a communicative encounter. It will only *re-frame* if the current frame can no longer be used or does no longer seem desirable in terms of expected utility.

4. LEARNING WITH FRAMES

To learn the long-term usefulness of different frames in different situations based on rewards obtained during previous conversations, we use the *hierarchical reinforcement learning* (HRL) framework of *options* [5] and re-interpret interaction frames as "macro-actions" [7]. The options framework reduces the overall number of states in a Markov decision process (MDP) by combining the available actions into so-called options that can be applied over several decision steps and (hopefully) optimally solve sub-problems of the original MDP. Semi-MDP (i.e. state history dependent) variants of learning methods such as Q-learning can then be used to optimise the long-term "meta"-strategy over these macro-policies.

To allow for the acquisition and adaptation of a set of frames from actual conversations, [3] further describes a method which views frames as clusters in the space of interactions and aims at maximising the quality of the overall clustering.

Combining these learning mechanisms with expected utility maximisation among the choices still offered by the currently selected frame (since variables provide certain degrees of freedom) yields a two-layer hierarchical model for communicative decision making.

5. EVALUATION

The adequacy of our approach has been tested in the Link Exchange Simulation system LIESON [4]. In this system, agents represent web site owners who hold different views of the contents of other web sites in terms of numerical *private ratings*. At the same time, they can express their opinion about others sites by laying numerically weighted links toward these, which then function as a kind of *public ratings*. The primary goal of agents in LIESON is to increase the dissemination of their own opinion through appropriate linkage structures, and for this purpose they negotiate with each other over mutually beneficial linkage.

Experiments with both simple *proposal-based negotiation* and more complex *argumentation-based negotiation* frames prove that m^{2}InFFrA manages to integrate frame application and learning capabilities in a coherent social reasoning architecture that is able to operate successfully *in conjunction with* other (sub-social) agent activity in open MAS. See [6, 2, 7] for details and specific results.

6. CONCLUSION

In this paper, we have provided a brief summary of an integrated framework for reasoning about conversation patterns in multiagent systems. To our knowledge, it is the first architecture for reasoning about interaction patterns that combines boundedly rational decision-making based on a probabilistic interpretation of agent communication processes with hierarchical reinforcement learning for the long-term optimisation of communication strategies.

7. REFERENCES

[1] F. Dignum, editor. *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Artificial Intelligence*. Springer, Berlin, Germany, 2004.

[2] F. Fischer and M. Rovatsos. Reasoning about Communication: A Practical Approach based on Empirical Semantics. In *Proceedings of the 8th International Workshop on Cooperative Information Agents (CIA)*, Erfurt, Germany, 2004. Springer.

[3] F. Fischer, M. Rovatsos, and G. Weiß. Acquiring and Adapting Probabilistic Models of Agent Conversation. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Utrecht, The Netherlands, 2005. ACM Press.

[4] LIESON. http://www7.in.tum.de/∼rovatsos/lieson, 2002-2004.

[5] D. Precup. *Temporal Abstraction in Reinforcement Learning*. PhD thesis, Department of Computer Science, University of Massachusetts, Amherst, 2000.

[6] M. Rovatsos. *Computational Interaction Frames*. PhD thesis, Department of Informatics, Technical University of Munich, 2004.

[7] M. Rovatsos, F. Fischer, and G. Weiss. Hierarchical Reinforcement Learning for Communicating Agents. In *Proceedings of the Second European Workshop on Multiagent Systems (EUMAS)*, Barcelona, Spain, 2004.

[8] M. Rovatsos, M. Nickles, and G. Weiß. Interaction is Meaning: A New Model for Communication in Open Systems. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Melbourne, Australia, 2003.