Learning Policies through Argumentation-derived Evidence

(Extended Abstract)

Chukwuemeka D. Emele
University of Aberdeen,
Aberdeen, AB24 3UE, UK
c.emele@abdn.ac.uk

Timothy J. Norman
University of Aberdeen,
Aberdeen, AB24 3UE, UK
t.j.norman@abdn.ac.uk

Frank Guerin
University of Aberdeen,
Aberdeen, AB24 3UE, UK
f.guerin@abdn.ac.uk

Simon Parsons
Brooklyn College, City
University of New York, USA
parsons@sci.brooklyn.cuny.edu

ABSTRACT

We present an efficient approach for identifying, learning and modeling the policies of others during collaborative activities. In a set of experiments, we demonstrate that more accurate models of others' policies (or norms) can be developed more rapidly using various forms of evidence from argumentation-based dialogue.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Algorithms, Experimentation

Keywords

Argumentation, Machine learning, Policies, Norms

1. INTRODUCTION

Many distributed problem solving scenarios require the formation of a team of collaborating agents. In such scenarios agents often operate under constraints, which determine to a large extent, how they go about their assigned tasks. When these constraints are part of the standard operating procedures of the agents or the organisations they represent, we refer to them as *policies* (also known as norms). Although agents may have prior assumptions about the policies that constrain the activities of others, these models are often incomplete and may be inaccurate. In a set of experiments, we investigate how information acquired through dialogue, which we call argumentation-derived evidence (ADE), can be effectively exploited to learn better models of agents' policies. Our claim is that, through the use of argumentation-derived evidence, better (more complete and correct) models of the policies of others can be learned. We demonstrate the validity of this claim through experiments conducted in a resource provisioning scenario [1]. The research presented in this extended abstract represents the first model for using evidence derived from argumentation to learn underlying social characteristics (e.g. policies/norms) of others. There is, however, some prior research in combining machine learning and argumentation, and in using argument structures for machine learning. Rovatsos et al. [4] use hierarchical reinforcement learning in modifying symbolic constructs that regulate agent *conversation patterns*, and argue that their approach could improve an agent's conversation strategy. In other research, Možina et al. [2] propose a novel induction-based machine learning mechanism using argumentation. In our work, we used information-seeking dialogue [6] to obtain evidence from the interaction and learned the entire sequence as against a segment (frame) of the interaction [4].

Cite as: Learning Policies through Argumentation-derived Evidence (Extended Abstract), Chukwuemeka D. Emele, Timothy J. Norman, Frank Guerin and Simon Parsons, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2. EXPERIMENTS

The scenario adopted in this research involves agents acting as resource *seekers*, that interact with potential *providers* to resource a plan. In order to gather useful evidence, agents engage in argumentative dialogue and attempt to learn from dialogical encounters with other agents. The enactment of both *seeker* and *provider* roles are governed by individual policies that regulate their actions. A *seeker* agent requires resources in order to carry out some assigned tasks. The *seeker* agent generates requests in accordance with its policies and negotiates with *provider* agents based on these constraints. On the other hand, *provider* agents have access to certain resources and may have policies that govern the provision of such resources to other members of the team. *Provider* agents do not have an unlimited pool of resources and so some resources may be temporarily unavailable. By a resource being available we mean that it is not committed to another task (or agent) at the time requested.

In this study we used three different machine learning mechanisms: *Decision tree learning*, *Instance-based learning* and *Rule-based learning*. These three mechanisms represent very different classes of machine learning algorithms. The rationale for exploring a range of learning techniques is to

Table 1: Average percentage of policies classified correctly and standard deviation

Tasks	C4.5 - ADE	C4.5 + ADE	k-NN - ADE	k-NN + ADE	Prism - ADE	Prism + ADE
1000	58.3 ± 15.1	60.3 ± 14.4	65.2 ± 9.8	71.1 ± 9.0	66.7 ± 8.2	67.7 ± 7.7
2000	69.2 ± 16.6	75.0 ± 12.6	71.0 ± 7.8	85.9 ± 7.3	71.7 ± 6.0	87.1 ± 6.4
3000	75.1 ± 12.0	83.6 ± 6.5	75.3 ± 5.3	92.0 ± 4.6	78.7 ± 8.4	94.1 ± 4.2
4000	82.1 ± 12.3	89.9 ± 5.2	80.7 ± 3.8	96.8 ± 3.1	84.3 ± 6.5	96.6 ± 4.1
5000	85.3 ± 8.9	93.0 ± 3.4	81.0 ± 4.1	97.3 ± 3.6	87.4 ± 6.0	97.5 ± 2.6
6000	88.2 ± 8.2	95.6 ± 5.1	82.0 ± 3.8	98.4 ± 1.7	90.6 ± 5.3	99.2 ± 1.0

demonstrate the utility of argumentation-derived evidence regardless of the machine learning technique employed. Thus, our hypothesis is that the use of evidence acquired through argumentation significantly improves the performance of machine learning in the development and refinement of models of other agents. *Decision tree learning (C4.5)*: In this approach, the policies are conceived as concepts of an agent. Agent policies are represented as a vector of attributes and these attributes are communicated back and forth during negotiation. The C4.5 algorithm is then used to classify each set of attributes (policy instance) into a class (that is, grant or deny). In future encounters, the *seeker* agent attempts to predict the policies of the provider based on the model it has built. 2. *Instance-based learning (k-Nearest Neighbours)*: In this approach, policies are stored as instances and new policy instances are classified based on the closest training examples in the feature space. 3. *Rule-based learning (Prism)*: This approach encode policies as rules, and use sequential covering to induce the rules by selecting attribute-value pairs that satisfy the rule- one rule at a time. The policies covered by the rule are then removed and the process iterates until all the rules are covered.

3. RESULTS

Experiments were conducted with *seeker* agents initialised with random models of the policies of providers. 100 runs were conducted for each case, and tasks were randomly created during each run from 375 possible configurations.

Table 1 illustrates the effectiveness of identifying and learning policies through argumentation-derived evidence using the three machine learning techniques described above. It shows the percentage of policies classified correctly and their standard deviations. In each case the model of others' policies is recomputed after each set of 1000 tasks. For all three machine learning techniques considered, the percentage of policies predicted correctly as a result of exploiting evidence derived from argumentation was consistently and significantly higher than those predicted without such evidence. This result shows that the exchange of arguments during practical dialogue enabled agents to learn and build more accurate models of other agents' policies much faster than scenarios where there was no exchange of arguments.

Tests of statistical significance were applied to the results. Using linear regression, the analysis of variance (ANOVA) shows that as the number of tasks increases, each of the three learning techniques (with or without argumentation-derived evidence) consistently converges with a 95% confidence interval. Furthermore, for all the pairwise comparisons, the scenarios where argumentation-derived evidence was combined with machine learning techniques consistently yielded higher rates of convergence ($p < 0.02$) than those without additional evidence. These results confirm our hypothesis.

4. CONCLUSIONS

The results of these experiments have shown that evidence derived from argumentation can have a statistically significant positive impact on identifying, learning and modeling others' policies during collaborative activities. Accurate policy models can inform strategies for advising human decision makers on how a plan may be resourced and who to talk to [5], and may aid in the development of more effective strategies for agents [3]. Our results demonstrate that significant improvements can be achieved by combining machine learning techniques with argumentation-derived evidence.

Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

5. REFERENCES

[1] C. D. Emele, T. J. Norman, F. Guerin, and S. Parsons. Learning policy constraints through dialogue. In *Proc. of the AAAI Fall Symposium on The Uses of Computational Argumentation*, pages 20–26, Virginia, USA, 2009.

[2] M. Možina, J. Žabkar, and I. Bratko. Argument based machine learning. *Artificial Intelligence*, 171(10-15):922–937, 2007.

[3] N. Oren, T. J. Norman, and A. Preece. Loose lips sink ships: A heuristic for argumentation. In *Proc. of the 3rd Int'l Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2006)*, pages 121–134, 2006.

[4] M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiss. Practical strategic reasoning and adaptation in rational argument-based negotiation. In *Argumentation in Multi-Agent Systems*, volume 4049 of *LNCS*, pages 122–137. Springer-Berlin, 2005.

[5] K. Sycara, T. J. Norman, J. A. Giampapa, M. J. Kollingbaum, C. Burnett, D. Masato, M. McCallum, and M. H. Strub. Agent support for policy-driven collaborative mission planning. *The Computer Journal*, page bxp061, 2009.

[6] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning.* SUNY Press, Albany, NY, USA, 1995.