Incorporating Social Practices in BDI Agent Systems*

Extended Abstract

Stephen Cranefield
University of Otago
Dunedin, New Zealand
stephen.cranefield@otago.ac.nz

Frank Dignum
Umeå University
Umeå, Sweden
frank.dignum@umu.se

KEYWORDS

Social practices; BDI agents; Jason; Metadeliberation

ACM Reference Format:
Stephen Cranefield and Frank Dignum. 2019. Incorporating Social Practices in BDI Agent Systems. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Imagine the scenario where an elderly person, living alone, is assisted by a care robot. The robot makes sure that the person gets up every morning and that he drinks some coffee and takes his morning pills (if needed). Then they read the newspaper together, which means that the person looks at the pictures in the paper and the robot reads the articles out loud.

When agents in the role of this type of personal assistant or care robot have to interact with humans over a longer time period and in a dynamic environment (that is not controlled by the agent), the interaction management becomes very difficult. When fixed protocols are used for the interaction they are often not appropriate in all situations and cause breakdowns and consequent loss of trust in the system. However, to have real-time deliberation about the best response during the interaction is not very scalable, because in real life the contexts are dynamic and complex and thus the agent would need to take many parameters into consideration at each step. Thus we need something in between a completely scripted interaction that is too brittle and a completely open interaction that is not scalable.

As we have done before in the agent community, we take inspiration from human interactions and the way they are managed by individuals. We classify situations into standard contexts in which a certain *social practice* can be applied. Social science has studied this phenomenon in social practice theory. Social practice theory comes forth from a variety of different sub-disciplines of social science. It started from philosophical sociology with proponents like Bourdieu [3] and Giddens [6]. Later on Reckwitz [9] and Shove et al. [11] have expanded on these ideas, and Schatzki [10] made some valuable contributions.

These authors claim that important features of human life should be understood in terms of organized constellations of interacting persons, which together constitute social practices. People are not just creating these practices, but our deliberations are also based on the fact that most of our life is shaped by social practices. Thus, we use social practices to categorize situations and decide upon ways of behaviour based on social practices.

Unfortunately social practice theory has not been widely used in computer science or in HCI and thus there are no ready-to-use tools in order to incorporate them in agents. It is clear from the above description that social practices are more than just a protocol or a frame to be used by the agent in its deliberation. Therefore, in this paper we make the following contributions. We propose a mechanism for BDI agents to maintain awareness about active social practices, and to leverage their existing plans to act in accordance with these practices. This takes the form of a set of metadeliberation plans, which can be directly executed by Jason [2] agents, or treated as a specification for an optimised implementation in an extended agent platform. These plans have been deployed in the (simulated) care robot scenario, to confirm that awareness of and adherence to a social practice enables the robot to have a more successful interaction with the patient over a longer period of time. As some of the features needed to implement this scenario, and to support our metadeliberation plans, are not currently available in Jason, we also developed a Jason metainterpreter, which provides this extended functionality, but can also be used independently to support other research on extensions to BDI practical reasoning.

2 SOCIAL PRACTICES FOR BDI AGENTS

Social practices have been conceptualised in the context of multi-agent systems [5] and formalised in dynamic logic [4]. Based on these accounts, we focus on these aspects of social practices: (a) they are relevant in specific contexts, defined in terms of the actors, resources and places involved; and (b) they are modelled as plan patterns, structured as a set of partially ordered landmarks, each with an associated *purpose* (a goal) and a sequence of actions that is a partial prescription for reaching the landmark. Thus, to follow social practices, a BDI agent must maintain awareness of the currently relevant practices, track the status of landmarks, and proactively use its plans to achieve the landmarks' purposes by selecting the plans that include the landmarks' prescribed actions.

3 THE CARE ROBOT SCENARIO

In this section we elaborate on the care robot scenario outlined in the introduction, and describe how we have modelled and implemented it using Jason [2].

We assume the high-level operation of the robot is based on a BDI interpreter, and that it comes equipped with goals and plans to trigger and enact its care activities (most likely with some customisation of key parameters possible). In this section, we consider

*The full paper can be found at http://arxiv.org/abs/1903.03189. For source code, see https://github.com/scranefield/jason-social-practices

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

only a small subset of the robot's duties: to wake the patient at a certain time in the morning, to provide coffee as required, and to provide mental stimulation. We do not specify any goals of the robot outside the practice, but normally the care robot would also have its own goals such as powering its battery, (vacuum) cleaning a room and taking care of the health of the patient.

To perform most effectively, the robot should choose, for a given context, the plans for each goal that will achieve the best outcomes for the patient, and furthermore, consider constraints on goal orderings that arise from context-specific preferences and habit. For example, if the patient prefers to be woken at a certain time in a given context (e.g. when his family is due to visit) and/or in a certain way (e.g. by the curtains being opened), his mood is likely to be adversely affected if he is woken at a different time. Optimising the patient's experience therefore involves context-specific coordination between the robot's and patient's activities. We model this coordination for our scenario as a social practice involving landmarks to have the patient awake, have his pills taken, to be served coffee, and to be mentally stimulated.

As social practices provide patterns of coordination for multiple agents in terms of landmark states rather than explicit sequences of actions, they do not make limiting assumptions about the temporal ordering of actions and when their effects occur. Only the landmarks themselves are explicitly temporally ordered. To illustrate this we include some temporal complexity in the scenario by including in our scenario a durative action (making coffee takes a period of time to complete), an action with a delayed effect (after opening the curtains, the morning light will eventually wake the patient), and a joint durative action (reading the newspaper together has its desired effect only if both agents perform it during overlapping time intervals).

4 IMPLEMENTATION

Maintaining awareness of social practices, and contributing to them in an appropriate way, requires agents to detect when each known social practice becomes active or inactive, to monitor the state of the landmarks in an active social practice, and to trigger the appropriate activity if a pending landmark has an action for the agent. This is a type of meta-level reasoning that the agent should perform periodically, and it may override the performance of any standard BDI processing of goals, which is not informed by social practices. In particular, the preferences of individual agents (in particular, in terms of plan selection) are overridden by an active social practice. We note that, on an abstract level, the same was done in [1] where the plan pattern was translated into a global pattern in Drools (a Java based rule engine) and the specific interactions within each phase were programmed in a chatbot.

The question then arises of how best to implement such a meta-level reasoner in a BDI architecture. The best performance can, no doubt, be achieved by extending a BDI platform using its underlying implementation language. However, this would require significant knowledge of the implementation and the use of an imperative coding style that is not best suited to reasoning about goals [7] and for rapid prototyping and dissemination of new reasoning techniques. Therefore, in this work we define the meta-level reasoner as a plan for a `metadeliberate` goal that reasons about social practices, sleeps and then calls itself recursively.

The plan determines which social practices have all their requirements satisfied, selects (currently) one, and updates a belief recording the selection. Any landmarks being monitored are then checked to see if their purpose has been fulfilled. If so, a belief about their completion is added. The plan then sleeps, before creating a goal to re-run itself.

A second plan handles a changed belief about the selected social practice. For each of its landmarks, if the agent already has an intention to achieve the landmark's purpose, that intention is suspended. A plan is also temporarily added to ensure that if some other active plan of the agent separately creates this intention, it will be immediately suspended. For each landmark in the social practice that has no prior landmarks, a goal is created to activate it.

Landmark activations are handled by a third plan. A belief recording that the landmark's purpose should be monitored is added, then the action (if any) associated with the landmark, for the current agent, is processed. First, an attempt is made to find a way to use the agent's plans to achieve the landmark's purpose, while ensuring that the specified action is performed. If such a solution is found, it is recorded as a path through a goal-plan tree [8] and passed to our Jason metainterpreter—a specialisation of Winikoff's AgentSpeak metainterpreter [12] for use with Jason, which is extended to accept goal-plan tree paths and to implement durative and joint actions. If no such solution is found, the action is performed directly.

Finally, a fourth plan handles completed landmarks—those for which the purpose has been achieved. Any suspended intentions for the purpose are succeeded, the belief stating that the landmark should be monitored is retracted, and the temporary plan mentioned above is removed. The plan then checks for subsequent landmarks that should now be activated (if all their prior landmarks are completed), and finally adds a belief that the social practice has completed if all its landmarks are completed.

5 CONCLUSIONS

We have argued that for interactive settings, as sketched in our scenario, the use of social practices is a good compromise between using a fixed interaction protocol and deliberation and planning from scratch at each point during the interaction. We proposed a mechanism for a BDI agent to maintain awareness about and contribute towards the completion of social practices, implemented as a metadeliberation plan for Jason agents. We also developed a Jason metainterpreter to support this plan and our care robot scenario.

Our approach allows BDI agents to use their existing plans to achieve social practice landmarks, with action choices in landmark specifications overriding the preferences of local plans. In future work we intend to investigate further ways in which social practices and local plans interact. We also intend to develop elaborate scenarios that use all aspects of a social practice, and compare these with agent implementations where no social practice is used, both in terms of the outcomes of the agent and the ease of design of the agents.

REFERENCES

[1] Agnese Augello, Manuel Gentile, and Frank Dignum. 2016. Social practices for social driven conversations in serious games. In *Proceedings of the 4th International Conference of the Games and Learning Alliance*. Lecture Notes in Computer Science, Vol. 9599. Springer, 100–110.

[2] Rafael H. Bordini, Jomi Fred Hübner, and Michael Wooldridge. 2007. *Programming multi-agent systems in AgentSpeak using Jason*. Wiley.

[3] Pierre Bourdieu (trans. R. Nice). 1972. *Outline of a theory of practice*. Cambridge University Press.

[4] Frank Dignum. 2018. Interactions as social practices: Towards a formalization. *arXiv* (2018). https://arxiv.org/abs/1809.08751

[5] Virginia Dignum and Frank Dignum. 2015. Contextualized planning using social practices. In *Coordination, Organizations, Institutions and Norms in Agent Systems X*. Lecture Notes in Computer Science, Vol. 9372. Springer, 36–52.

[6] Anthony Giddens. 1979. *Central problems in social theory: Action, structure and contradiction in social analysis*. University of California Press.

[7] Brian Logan. 2018. An agent programming manifesto. *International Journal of Agent-Oriented Software Engineering* 6, 2 (2018), 187–210.

[8] Brian Logan, John Thangarajah, and Neil Yorke-Smith. 2017. Progressing intention progression: A call for a goal-plan tree contest. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems*. ACM, 768–772.

[9] Andreas Reckwitz. 2002. Toward a theory of social practices. *European Journal of Social Theory* 5, 2 (2002), 243–263.

[10] Theodore R. Schatzki. 2012. A primer on practices. In *Practice-Based Education*. Practice, Education, Work and Society, Vol. 6. SensePublishers, Rotterdam, 13–26.

[11] Elizabeth Shove, Mika Pantzar, and Matt Watson. 2012. *The dynamics of social practice*. SAGE Publications.

[12] Michael Winikoff. 2006. An AgentSpeak meta-interpreter and its applications. In *Programming Multi-Agent Systems, Third International Workshop, ProMAS 2005*. Lecture Notes in Computer Science, Vol. 3862. Springer, 123–138.