Social Deliberation for Organizations

Frank Dignum [1] and **Virginia Dignum** [2] and **Rui Prada** [3] and **Catholijn Jonker** [4]

Abstract. Socially interconnected systems should be able to exhibit, understand, and reason about social behavior, in order to support lasting interconnections that show realistic and desirable properties. We claim that a fundamental new approach based on social concepts is needed to build agents functioning in these socially interconnected systems. In this paper, we explore how the concepts of social practices and social identities can be used to structure deliberations about interactions. We then show the consequences for the architecture and reasoning capabilities of these systems.
[5]

1 INTRODUCTION

Given that inter-organizational systems are increasingly more complex and open ended, using fixed protocols to guide and monitor interactions between participants is no longer viable. Instead, agents in these systems should be endowed with forms of social intelligence that allows for them to deliberate about socially realistic interactions. Therefore we propose the option in which the agents take some of the burden of this flexibility by endowing them with a social deliberation engrained in the core of their reasoning, such as have been developed by humans[3]. We start from the premise that social interaction (and reality) is not given, but socially constructed [13],[2]. Current systems still have very limited understanding of their context, and of their social role. They are therefore not able to reason about their identity and goals in a social context, and therefore cannot be expected to function outside the situations they've been designed for and actually create the social interactions as required.

Currently, sociability is engineered into the system, in a situational and context dependent way. Social signals are not appraised as such but implicitly built into their functionality. Therefore, the behaviour of the system is not conceived as social outside that particular context, and the system is not able to adapt to significant changes. This implies that reuse in different social contexts (such as interacting with different types of suppliers) or cultures (connecting with organizations abroad) often requires a complete re-engineering of the system. A next step forward in AI, is the ability to perceive, reason about and exhibit *social* intelligent behaviour. This will require a framework containing explicit social principles that can be described, represented and manipulated in a symbolic way.

[1] Utrecht University, The Netherlands, email: F.P.M.Dignum@uu.nl
[2] Delft University of Technology, The Netherlands, email: M.V.Dignum@tudelft.nl
[3] INESC-ID and Instituto Superior Tecnico, Universidade de Lisboa, Portugal, email: rui.prada@tecnico.ulisboa.pt
[4] Delft University of Technology, The Netherlands, email: C.M.Jonker@tudelft.nl
[5] IAT4SIS '14, August 18 2014, Prague, Czech Republic Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM 978-1-4503-2890-6/14/08...$15.00. http://dx.doi.org/10.1145/2655985.2655987

We argue that deliberative, social, and physical principles must be considered first class components of a computational theory of social intelligence. Being socially intelligent requires a keen understanding of the principles of social reality, and the ability to link social interpretations with individual goals into plans and vice versa. The interrelationship between social and physical contexts is such that the social context defines the possible social interpretations of the physical reality and limits the set of admissible actions; and the physical context determines and constrains the possible social contexts. For example, a raised hand can mean many things: in a class room: a question, in an auction hall: a bid, on the street: a greeting or threat.

It is clear that splitting the context in a social and physical one adds quite some complexity. If an agent plans for a social goal (such as gaining acceptance in a group of companies) it needs to plan physical actions to reach such a goal. Thus any plan has both social as well as physical consequences. The agent has to deliberate about all these elements in order to decide upon the best cause of action. It might be clear that agents need some structures in order to limit and facilitate this process. In this paper we describe the first steps towards agents which can be called socially intelligent.

In order to make our arguments more concrete we sketch a use case scenario in section 2 to illustrate our points. In section 3, we identify the key elements of a framework to build socially intelligent systems. We describe how motives can be used to give an agent both social as well as physical direction. Then we describe how social identity, norms, values and social practices can be used as means to preserve both consistency (at individual and society level) as well as simplicity. In section 4, we propose an abstract architecture for social reasoning built on top of the concepts introduced in section 3. In section 5, we sketch how this architecture can be used in order to determine a course of action in the example scenario, described in section 2, in an efficient way. In section 6 we discuss the social practice as a process and how this process determines its effectiveness. Finally, we draw some conclusions and indicate future work in section 7.

2 SCENARIO

Consider the situation where a company A is deliberating on their response to a merging bid from company P. The companies are competitors on the pharmaceutical market. However, the bidding company. P, is a large, profit-driven, company. Company A is much smaller, and geared towards research and development. Accepting the bid might give A many new possibilities to market their products. However, A might also loose flexibility to do blue sky research and might loose jobs that way. Members of company A act according to four core values: innovation, disregard for money (profit), independence and intellectual superiority.

In a meeting between managers of both companies to discuss the

merger, their differences became apparent. Company P managers wear a suit and tie, while managers of company A are casually dressed. Moreover, the managers of company A use the meeting as an opportunity to express their "independence" and "intellectual superiority". They act aggressively, being impolite, interrupting, and insulting members of the other group. As a consequence, negotiations reach an impasse, and the government decides to intervene. Given the economic and labour consequences of cancelling the merger, the government has a large stake in the process.

However, the presence of the government in the negotiations may lead to different results. On the one hand, both companies may perceive the government officials as meddling or repressing, and united by their common identity as captains of industry, combine together to oppose the government as it limits their independence. On the other hand, the intervention by the government may lead the managers to realize that they have common responsibility towards the country and their employees, and will increase their efforts to reach a suitable merger agreement.

3 SOCIAL FRAMEWORK

In order to create truly social intelligent systems, we need to start from an understanding of the motives for behaviour [6]. Often, human behavior results from a need to balance between novelty and control I.e. in the one hand, seeking out new situations, while, on the other hand, trying to avoid uncertainty and thus strive to control the environment. The balance between the two forces is different between persons, but prominent in the way people deal with social interactions. However, it does not readily indicate how people choose concrete actions. Thus we look at theories of human motivations [10] in order to get more concrete handles on drivers of behaviour. McClelland argues that there are a number of basic natural incentives that give rise to some motives. Besides the biological (homeostatic) motives such as hunger and need for sleep (which are, in fact, not very salient in most of the social situations), McClelland distinguishes four motives: (1) *achievement*, (2) *power*, (3) *affiliation* and (4) *avoidance*.

Achievement Leads to explorative behaviour and satisfies the need for novelty.

Power is about trying to have an impact on the world, and includes both the need to control both its own and other's actions.

Affiliation drives people to seek the company of others and to establish and maintain positive interactions with others.

Avoidance leads to self preservation, seeking certainty, and emotional regulation, and fosters the categorization and simplification of behaviour so that it becomes more standardized (and thus predictable).

In the next subsections we discuss a number of mechanisms that play an important role in human deliberation and that can be used for social intelligent agents as well. These mechanisms represent concrete ways to fulfil motives at a very high level of abstraction. We describe three constructs that are used in the Social Sciences to ensure consistency of behaviour over time: Identity, Norms and Routines/Habits. Finally, we discuss the concept of *Social Practices* as a way to describe the combination of social and physical context with respect to a (standard) course of action.

3.1 Identity and Values

A social intelligent system must be able to perceive itself and (its position in) the social world. People position themselves, and others, in terms of membership of, possibly many, social groups (i.e. reference groups) and social goals are often based on comparison with others [14] [15]. For example, if you want to be a good CEO this means that you identify yourself (at least partly) as a "CEO" and you need to know the position and activities of some (prototypical/ideal) good CEOs such that you can ascertain what kind of action is needed to become respected in that group. Some reference groups are quite stable, such as family and profession while others are more volatile, such as the group of people in a shop or at a meeting. People have different emotional attachments to each of the social groups, which elicits social goals to maintain and pursue certain identities.

What constitutes a "good CEO" relates to the set of values and their priorities associated with a reference group. Many definitions of values exist and many research communities use them in different ways. We see them as criteria with which pairs of situations can be ordered. E.g. the value "environmental friendly" can be used to compare two situations on the basis of how well the nature is preserved in each of them. It can very well be that another value, such as, "comfort" will sort the two situations exactly the other way around. Values are used to reconcile the different reference groups the person belongs to, such that her behaviour is consistent (and expected).

From the above it follows that once a situation triggers an identity in the agent (or if it selects it proactively), the identity comes with a set of values and thus prescribes a certain type of behaviour. For example, when a person has the identity of a research manager she will support researchers in performing their experiments at any time of the day. There is no deliberation about whether she would prefer to go shopping or go on a date at that moment. Identities are also social, because they give people the sense of belonging to the reference group. The identity is visible for others and expectations can be formed on the basis of knowledge about the reference group. Thus, the identity provides consistency of behavior on an individual level (because an individual with an identity will behave according to that identity) as well as on social level (because all individuals with a certain identity will behave similarly in the same context).

3.2 Norms

Norms are the second construct that can be used to categorize and classify behaviours. Norms specify behaviours that promote values. We will not get into all the different kinds of norms and description of their properties at this place (we refer to [1] for a recent overview on norms in multi-agent systems). For the purpose of this paper the function of norms as behavior regulation is the most important.

Norms also have an individual and social side just like identities. When an individual has accepted a norm it means that that individual will act according to that norm (in the appropriate context where the norm is active). Thus the norm ensures individual consistency of behaviour. The norm also has a social side, because norms indicate what is socially acceptable behaviour. Thus, they ensure consistency of behaviour, not just over one individual, but over all persons for which the norm is active. Often the activation conditions of norms refer to reference groups that are also used for identities. It, therefore, seems logical to include the norms pertaining to particular reference groups with the description of that identity. E.g. the "researcher" identity will come with some norms on how to behave at work and meetings. It should be kept in mind though that we do not assume a particular order of importance between norms and identities. The particular context will determine the most salient aspect and based on that the other aspects can be connected and become salient as well.

3.3 Habits

According to [11] habits are psychological dispositions to repeat past behaviour. Our daily life is full of habits, as ways to simplify decision making. When a certain behaviour has become a habit we do not deliberate about that behaviour anymore, but just repeat it whenever it is triggered by the context. According to the psychology literature habits are related to goal directed behavior and their origin can often be traced to a deliberate action with a particular goal.

As many behaviour is performed out of habit, habits can also be used to predict behaviour. This is explicitly done in the Consumat agent model that is used for consumer simulations [8]. It makes the agents simple and the focus can be on the cases where habits are broken. Habits also lead to consistent behaviour. Exactly because they are repetition of the same behaviour in similar conditions.

3.4 Social Practices

Having looked at several mechanisms to simplify deliberations about actions while keeping consistency in the previous sections we now turn to the issue on how to manage the three elements of social intelligence mentioned in the introduction. We need to look for structured theories that govern the rules with which these three elements are related. In the social science in recent years the concept of *social practice* seems to target exactly the same problem. Researchers in social science [7, 12] have identified three broad categories of elements of social practices:

- Material: covers all physical aspects of the performance of a practice, including the human body and the actions that can be performed as part of the social practice. This relates to our physical aspects of a situation.
- Meaning: refers to the issues which are considered to be (socially) relevant with respect to that material, i.e. understandings, beliefs and emotions. This relates to our social aspects of a situation.
- Competence: refers to skills and knowledge which are required to perform the practice. This relates to our notion of deliberation about a situation.

One could see a social practice as an elaborate condition-action rule. First a situation is assessed in order to check whether a social practice is relevant. When it is relevant all material elements get a social interpretation (partly) determined by the social practice. Thus, a group of people can be given meaning as being "my friends", "a rival management team" or "managers", depending on the social practice that is activated. Within a social practice several behaviours are possible. Depending on the competence that a person has, she will disregard some. Then an action is chosen for which a kind of standard social effect is also expected within this social practice. Of course all elements from the previous sections play a role in social practices as well. People in a social practice are given social meanings through their social identities. The actions available in social practices are often (social) norms and when a social practice is often followed the actions within it become habits. In this way all the concepts of the social framework are directly related.

The components are linked by individuals when carrying out a practice. Each individual embeds and evolves meaning and competence, and adopts material according to his or her motives, identities, capabilities, emotions, and so forth, such that a practice can then be implemented as a composition of components. However, because the social practice is also a shared notion one can expect the other participants in the social practice to share the same interpretation of the situation. There is an (expected) shared situation awareness.

Individuals and societies typically evolve a collection of practices over time that can be adopted in different situations. Social practices are like social norms in that they emerge from individuals, but are not dependent on the individuals any more. They are continuously shaped when they are followed and can differ for individuals with different experiences. E.g. we all share an understanding of the greeting practice, but the exact behaviours and social connotations may differ. Moreover, depending on the situation, the personality and the skills of an individual, carrying out a practice will be a more automatic or a more deliberated process. This corresponds to the different modes of thinking proposed by psychologists[6] [9]:

- System 1, or fast thinking, operates automatically and quickly, with little or no effort and no sense of voluntary control. This includes recognition, perception, and orientation.
- System 2, or slow thinking, allocates attention to activities that demand a high amount of mental effort. Such activities include complex computations, rule following, comparisons, and weighing of options. The operations of System 2 are often associated with the subjective experience of agency, choice, and concentration.

When a social practice is experienced very often the interpretation of the situation can be done in a standard way and leads to a quick decision on an optimal behaviour. E.g. when driving a car we hardly ever think about using the shift or the clutch when changing gear. However, when learning to drive we have to first learn which are the salient elements in the environment that trigger the gear change. Thus, we react slower and need more attention for the driving behaviours themselves. This is an important issue for persons in crisis situations, where decisions have to be taken quick. Experienced persons will very quick distinguish the salient elements in the situation and decide which social practice is most salient and act according to it. In the next section we will go more in-depth into the architecture for individual deliberation based on the use of social practices.

4 ARCHITECTURE

Social practice theories are, until now, mostly descriptive. Recently, a model was proposed to investigate the emergence of social practices [7]. However, this model does not explain how social practices are used in deliberation and how they are influenced by and influence individual agents. The abstract architecture, depicted in Figure 1, is a first attempt to capture this interaction.

The architecture shows the prominent place of social practices as an early input into the deliberation process. Although the sensing still is the start of the deliberation it is actively guided by both the motives as well as the social practices. As soon as social practices are selected as potentially fitting with the current situation they will drive the search for salient features in the environment that fit that social practice as well and might be used for determining (further) actions. E.g. if a person is getting to a meeting it will look for extended hands. This does not mean that this is the only feature that is searched for! Several patterns can be searched in parallel. However, their number is limited to patterns that can be *expected* within a current social practice.

A similar drive to search for patterns comes from the motives. E.g when entering a meeting with many people a person with a high affiliation motive starts looking for (possibly known) persons to interact with, while a person with a low affiliation motive might scan the room for a place to stand quietly. Thus, we see that the parallel tracks

[6] The labels of System 1 and System 2 are widely used in psychology.

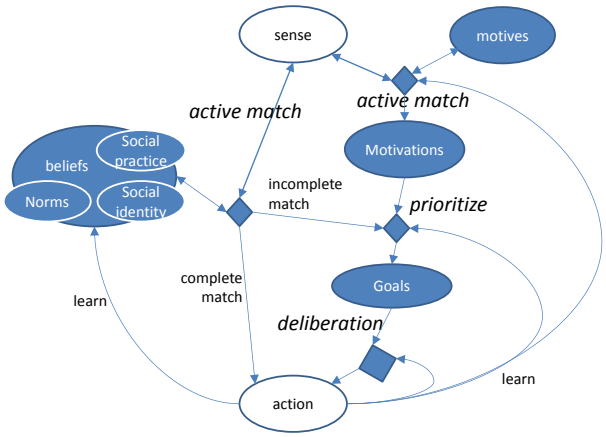

Figure 1. Abstract architecture for social reasoning

of pro-active and reactive behavior already start with the sensing behavior. Whether the motivation actually leads to setting a new goal is influenced by the social practices again. If the person is entering the room when a meeting is about to start, he might find a quiet place and people will be happy that he does not interfere with the meeting procedure. However, when there are after meeting drinks the social practice might dictate that the person mingles and talks with others.

If the course of action is not directly clear from the situation more deliberation takes place. In the figure the deliberation is represented by a simple kite symbol. This deliberation can contain a complex process itself, such as, the ones used in Fatima [4] or BRIDGE [5] containing emotions, goals, intentions, beliefs, roles, identity, etc. Suppose a CEO is strongly achievement driven. During a management meeting he will try to use a social practice that serves that motivation and maybe directs all discussion towards quick and efficient decision making. However, in order to achieve the ultimate goal of the decisions he needs the strong commitment of his staff. Thus, in order to get quick decisions with a high amount of commitment he will give opportunity for all persons to give their opinions before making a decision that reflects the inputs. Thus he will balance different aspects in choosing a course of action within a social practice.

When discussing the architecture on the basis of the above scenarios the social practices take a leading role in organizing possible courses of action. Note that our architecture does not depend on a fixed set of plans per goal nor that it needs a large set of plans to be searched through. The social practices combine material and social aspects in such a way that one can start from either side and check the appropriateness of the other aspect for the current situation. This avoids having to reason separately about both aspects and combining them afterwards. Having the social practices can also instantiate elements in the deliberation even if they are not totally clear from the initial interpretation of the context, such as the roles and expected goals.

The final aspect that we included in the architecture is the learning that takes place after the action has been executed. After each action the system should not just check whether the action succeeded or failed, but also whether it can use the result as feedback on the choices it made during the deliberation and whether it should refine or adjust its library of social practices. E.g. it might notice that it expects a handshake in the greeting practice while not everyone is shaking its hand. Thus it can extend this social practice with some alternative ways of greeting like bowing or hand waving. However, it might also learn that it successfully completed the social practice of decision making and update the salience of the plans it executed for this social practice. In this way it can update its memory even without explicitly storing every interaction. Finally, we should remark that where physical effects of actions can usually be measured with sensors, the social effects are often not visible and have to be derived from consequent actions of the partners. Thus, more subtle sensing and interpretation is needed to learn the most efficient social interaction patterns.

5 SOCIAL PRACTICES IN PRACTICE

The previous section illustrates how our proposed architecture for social intelligent systems could function. This section describes how the elements of identity, norm and social practices can be used in the scenario sketched in section 2 in order to provide efficient and consistent behavior, while still reacting to events in the environment.

People use their identities and motives, and their assessment of the current situation, to determine which action to take and what the meaning of the behavior of someone else is. Roughly, the deliberation process goes as follows (illustrated in Figure 2):

- Determine the most salient Identity for the situation
- Identify abstract Social Practices associated with that Identity
- Match the context to those Social Practices
- Identify a possible concrete instantiation for a matching Social practice
- Decide which action to take from the actions associated with the Social Practice (this choice depends on own goals, past experiences, skills, etc.)
- Take action
- Evaluate the result of actions as feedback to Identity (change of status and commitment) and to Social Practice (reinforcement)

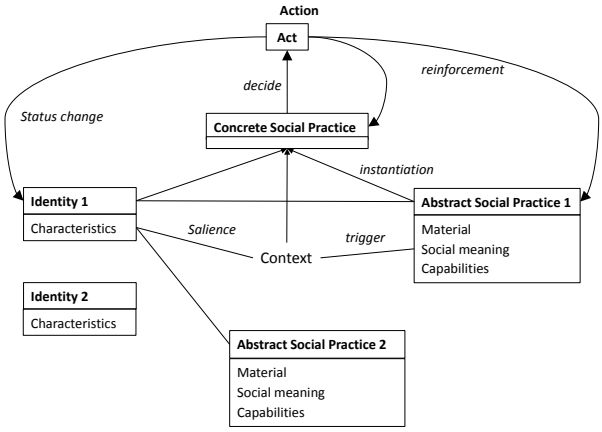

Figure 2. Abstract architecture for social reasoning

We show the use of this process by applying it to the deliberation by actors in our scenario. We assume three identities "Manager", "Government" and "Employee (of A)". All actors have the default "Citizen" identity. Figure 3 shows that an actor may decide that its identity of "employee of A" is the most salient in the situation where two groups of different companies encounter each other.

The situation triggers the instantiation of the social practice "Oppose out-group" to the social practice "Be aggressive towards managers P". The actor may then decide to take the action "Intellectually insult managers P". The result of this action feeds back into its identity (e.g. causing the meeting to result in an adverse result leads to a loss of status, which can lead to abandonment of the identity) and into its appreciation of the applicability of social practice.

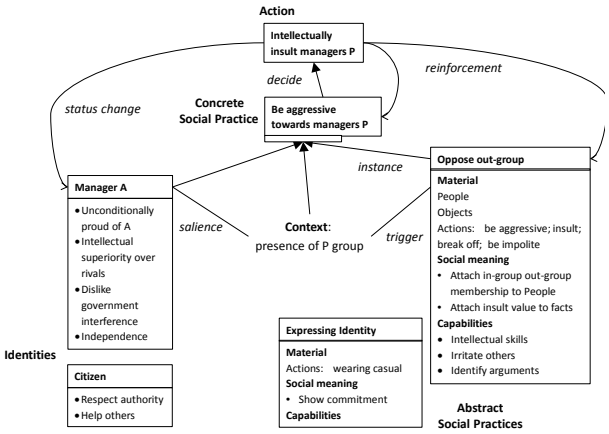

Figure 3. Social deliberation for a manager

In order to determine which of the actions associated with the social practice to take, the agent will take into account its goals, norms and habits, as described in section 4. In fact, the 'decide' activity identified in Figure 2 is achieved by the deliberation architecture illustrated in Figure 1.

In common situations, e.g., daily activities, agents will resort to habits, which are practices that require little to no deliberation and yield acceptable results (fast thinking). In such a case, there is a complete match between the information sensed from the context and some social practice the agent knows about (c.f. deliberation architecture in Figure 1). For instance, meetings with managers of a supplier company lead to standard greetings and exchanges of information plus possible new contract.

In other cases, deliberation is less direct and takes into account goals and motivations of the agent. The merger situation happens seldom and a manager might have worked very hard for many years to get his current position. Thus one manager might include the avoidance and power motive and the goal of preserving salary and position. However, the founder of the company may actively seek a fight to preserve the research values of the company. The reasoning here is that intellectual superiority over rivals is a social norm in the group and by upholding that norm and showing skill (capability) in that activity improves the view others have on him.

A similar deliberation mechanism for a government official is illustrated in figure 4. Getting called into the negotiations between the companies he automatically will look for consensus and tries to preserve the jobs of laborers as much as possible. That is, the match between sensed input, identity (of government official) and social practice is direct and complete, and little explicit deliberation is needed. As described in the scenario, this action may cause (unexpected) reactions from the managers, from both A and P. In this case the managers see the government official as a person limiting their independence, and unite to oppose the official. As a result the official now finds himself in an untenable situation. By taking into

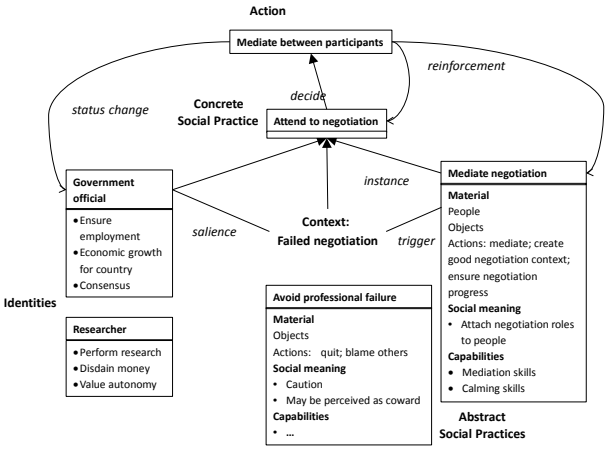

Figure 4. Social deliberation for a government official

account his own goals (e.g. negotiation success) and motives, this changes the context and triggers another social practice for the official, i.e.,"avoid professional failure", which leads him to quit the negotiations and blame the companies for it.

Another example is the case that the official is himself a former researcher. In this case, he may take its affiliation motive into account, that makes him wanting to express its belonging to the group of researchers, and decide to favor company A in the negotiations.

These examples show that many different behaviours can be expressed by the integration of the deliberation architecture described in section 4 and the social reasoning architecture depicted in Figure 2. Even though more extensive evaluations are needed, this scenario illustrates the applicability of the main concepts.

6 SOCIAL PRACTICES AS PROCESS

The previous section shows how existing social practices can be used to support individual deliberation about actions. However, that perspective does not highlight the social process of using social practices. When a social practice is used it always is done in a social context and thus has other individuals involved. The consequence is that when a social practice is successfully executed all the parties involved will get a positive feedback. This fact has as consequence that these parties will use the same social practice again (in the same context). Thus the fact that the social practice is successfully used leads to a high chance that it will be re-used. This, by itself, makes the social practice already more successful in the coordination process. I.e. if you can expect in a certain context that all parties involved will use a similar social practice this will facilitate the effective coordination. It right away shows the main value of standardization in business processes. It is not so much the exact format of the standard, but the fact that a standard is used which makes it effective.

The spread of a social practice depends on a number of factors. First of all, of course, how successful it has been to support a coordination. There is a balance here between generality of the practice and effectiveness of the support. If a social practice is applicable in many situations and contexts then inevitably the exact action patterns are also more general. If the generality of the actions leads to more miscoordination or coordination effort (to determine the exact actions to be performed) this might lead the social practice to be perceived as less successful. However, if the social practice is very specific and

only applicable in very specific situations as well it will not be used often (because it usually does not fit the context). This observation leads to a natural requirement to structure the social practices in a kind of hierarchies from general to very specific, such that individuals can easily move around the hierarchy in order to find the most specific level that is applicable in a context and that gives the most specific action support. If this practice is successfully executed not only that specific practice gets a positive feedback, but also the more general practices, thus reinforcing that whole branch.

A concern that is directly linked to the point above is the number of social practices that is applicable in a certain context. If only one practice is applicable it is easy for an individual to choose, but if this can only be achieved with very specific practices that are hardly ever applicable the maintenance and navigation of the social practice structure becomes very cumbersome. If the applicable contexts of social practices overlap a lot, the decision which social practice to choose might become difficult. This in itself would also lead to possible mismatches when different parties involved in the coordination in that context choose different social practices. Ideally we would have a unique social practice for each context. By linking them through generalization/specialization links along the different aspects (material/meaning/competence) it is then easy to navigate to the most specific uniquely applicable practice. Feedback on its performance should then again be given not only to the particular practice, but also (in lesser extent) to the practices linked to it.

Of course, in practice social practices can emerge in parallel. Whether one will dominate the others depends on the speed with which it spreads. This depends on how often parties involved in the practice participate in coordinations with different parties. The better a party is connected the more influence it gets in spreading a social practice. Also when a social practice is successful in many contexts it can be used more often and thus will spread quicker.

Finally the above considerations depend a lot on the context in which the coordination takes place. If the environment is very static and regular, one or two coordination mechanisms suffice to provide effective coordination. Thus few social practices will emerge and they will spread quick and converge quick. If the environment leads to many distinct contexts with their own specific requirements for coordination (like in human society) many more social practices may arise and also disappear again (when no longer useful). On an abstract level the process of convergence of social practices is similar to that of the convergence of terminology of ontologies in multi-agent systems. This also depends on (regular) interactions, the usefulness of the terms and the hierarchical shape of the ontologies. Some experimental results that can be used as examples for further simulations with social practices are described in [16].

7 CONCLUSIONS

We have shown that intelligent systems will need to be more socially aware of their context in order to take the appropriate action. As a result the deliberation of these systems should keep track of both its social and physical context when it deliberates about actions. If one would add the social context as an extra module (or aspect) to traditional intelligent systems this would overload the deliberation cycle and render it inefficient to cope with real time situations.

As a solution we have indicated several elements that can support the social behavior and that lead to a new type of architecture. Social identities, norms and habits are used as mechanisms to prioritize potential behaviours. They also serve to keep individual behaviour consistent over time. Norms and identities also have a strong social component in that they are shared by individuals in a society. Therefore, they can also be used to predict behaviour of others in known situations. The use of social practices facilitates the combination of social and physical aspects of a situation. Although these social practices might over constrain the combinations they lead to a good and natural reduction of the complexity.

This paper presents a first step towards truly social intelligent systems. There are many issues that are interesting for future research, such as dealing with the case that several social practices match a given situation, and emergence and evolution of social practices.

Acknowledgement
F. Dignum and Prada were supported by The Netherlands Institute for Advanced Study in the Humanities and Social Sciences (NIAS). Prada was also supported by Portugal's Fundação para a Ciência e a Tecnologia, under project PEst-OE/EEI/LA0021/2013.

REFERENCES

[1] Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leendert W. N. van der Torre. *Normative Multi-Agent Systems*. Dagstuhl Follow Ups, Germany, 2013.
[2] P.L. Berger and T. Luckmann. *The social construction of reality: a treatise in the sociology of knowledge*. Penguin Books, New York, 1996.
[3] C.L. Breazeal. *Designing Sociable Robots*. MIT Press, Cambridge, USA, 2004.
[4] J. Dias, S. Mascarenhas, and A. Paiva. Fatima modular: Towards an agent architecture with a generic appraisal framework. In *Proceedings of the International Workshop on Standards for Emotion Modeling*, 2011.
[5] F. Dignum, V. Dignum, and C.M. Jonker. Towards agents for policy making. In *MABS IX*, pages 141–153. Springer, 2009.
[6] F. Dignum, R. Prada, and G.J. Hofstede. From autistic to social agents. In *AAMAS 2014*, May 2014.
[7] G. Holtz. Generating social practices. *JASSS*, 17(1):17, 2014.
[8] M.A. Janssen and W. Jager. Fashions, habits and changing preferences: Simulation of psychological factors affecting market dynamics. *Journal of Economic Psychology*, 22:745–772, 2001.
[9] D. Kahneman. *Thinking, fast and slow*. Farrar, Straus & Giroux, 2011.
[10] D.C. McClelland. *Human Motivation*. Cambridge Univ. Press, 1987.
[11] David T Neal, Wendy Wood, Jennifer S Labrecque, and Phillippa Lally. How do habits guide behavior? perceived and actual triggers of habits in daily life. *Journal of Experimental Social Psychology*, 48(2):492–498, 2012.
[12] A. Reckwitz. Toward a theory of social practices. *European Journal of Social Theory*, 5(2):243–263, 2002.
[13] J.R. Searle. *The construction of social reality*. Penguin, 1996.
[14] H. Tajfel. Social identity and intergroup behavior. *Social Science Information*, 13(2):65–93, 1974.
[15] J. Turner, M. Hogg, P. Oakes, S. Reicher, and M. Wetherell. *Rediscovering the Social Group: A Self-Categorization Theory*. New York: Basil Blackwell, 1987.
[16] J. van Diggelen, R.J. Beun, F. Dignum, R.M. van Eijk, and J.-J.Ch. Meyer. Ontology negotiation in heterogeneous multi-agent systems: the anemone system. *Applied Ontology, An Interdisciplinary Journal of Ontological Analysis and Conceptual Modeling*, 2(3-4):267–303, 2007.