

Contextualized Planning Using Social Practices

Virginia Dignum¹ and Frank Dignum²

¹ Delft University of Technology, Delft, The Netherlands

M.V.Dignum@tudelft.nl

² Utrecht University, Utrecht, The Netherlands

F.P.M.Dignum@uu.nl

Abstract. Intelligent agents increasingly need to be aware of the social aspects of their context in order to take the appropriate action. However, existing techniques and platforms only provide partial solutions for this problem which do not take into account the full consequences of the social context. In this paper we propose to use ideas from social practice theory to support reasoning about action and planning in a social context.

We argue that putting social practices at the heart of the deliberation rather than use them as yet another aspect to be taken care of in the practical planning allows for more efficient planning. We provide a sketch of how this architecture provides some structure in the complexity of the deliberation process and balances between pro-active and reactive behaviour. The approach is demonstrated in a scenario taken from emergency management.

1 Introduction

Understanding the social contexts in which actions and interactions take place is of utmost importance for planning one's goals and activities. A system is context-aware if it can extract, interpret and use information about its context to adapt its plans to its current situation. Whereas people are pre-eminently able to understand context, computer systems are notorious for their inability to do so in general.

Social context is defined as the immediate physical and social setting in which something happens or develops. It includes the culture of an individual, and the people and institutions with whom they interact [1]. Within the agent community and in particular the COIN community the social context is often seen as consisting of the norms, organizations and institutions in which an agent is embedded. Subsequently it is investigated how these structures determine the behaviour of agents in application domains. E.g. in [21] we have shown the complexity of adding norms to BDI agents. Similarly in [2] it is shown how organization aware agents can be designed on the basis of BDI agents. From these and similar approaches it becomes clear that adding social context to BDI agents can complicate the deliberation process and often makes the standard deliberation inefficient (especially when used in real-time social environments).

In this paper, we argue that if we see social contexts not as an extra element to be added to an existing deliberation process, but rather as the *foundation* from which the deliberation process starts, it is possible to create more efficient deliberation processes. In a limited way, this has been shown for the case of norms in [15]. There the benefits were shown of using norms in the planning process rather than using them as a filter afterwards. In the current paper we take a broader perspective and include not only norms but all aspects of social context. We furthermore argue that *social practices* can be used to describe this social context in an efficient way.

Social practices give a means to choose between reactive behaviour in standard circumstances and pro-active behaviour that is necessary for social intelligent behaviour.

Existing agent platforms do not give much support to find this balance between pro-active and reactive behaviour. On the one hand, BDI agents, are primarily pro-active, goal driven. BDI implementations such as 2APL [6] or Jason [3], are particularly suitable to identify possible plans given a goal. Thus, for example when more than one plan is possible, 2APL agents are not able to identify which is the most suitable in a given context (the first plan that applies is followed), which may lead to plan revisions down the road, and as such, less efficient performance.

On the other extreme, we have completely reactive agents and Case-based Reasoning systems. These systems identify one possible action (or plan) given a situation by comparing that situation to the rules or known cases. Although this seems similar to the use of context it is limited in the sense that rules or cases are complete descriptions of situations. Only if all parameters are known or estimated can the action be determined. It depends also on the rule or case base whether “similar” cases can be found in order to derive a solution for the situation at hand. Usually this works well in a limited domain, but fails when the domains become dynamic or complex and cases are too scattered to provide answers for most situations.

Finally, Work Practice Modelling has been proposed to support the analysis of complex human-system interactions [19]. The rationale here is that understanding interactions requires going beyond formal procedures and information flows to analyse how people interact with each other. This approach is similar to the one we propose in this paper, but, as is the case in Case-Based Reasoning, current implementations are based on frames, and require the complete filling up of frame slots with situation information, and can thus not be used with incomplete information.

In this paper, we claim that agent plan generation can be enhanced by applying ideas from social practice theory. Social practices can be seen as a middle ground that combines the advantages of goal-directed and Case-Based Reasoning processes, by using social practices as heuristics to guide context-oriented plan identification. Social practice theory seeks to determine the link between practice and context within social situations [20]. Social practices refer to everyday practices and the way these are typically and habitually performed in (much

of) a society. Such practices as “going to work”, “meeting”, or “greeting” are routinely performed and integrate different types of elements, such as bodily and mental activities, material artefacts, knowledge, emotions, skills, and so on [16]. Social practices are similar for groups of individuals at different points of time and location. As such, they can be seen as ways to act in context, i.e. once a suitable practice is identified, people will use it as a ‘short cut’ to determine an action which does not require elaborate reasoning about the plan to follow. However, social practices are not just mere scripts in the sense of [14]. They support, rather than restrict deliberation about behaviour. E.g. the social practice of “going to work” incorporates usual means of transport that can be used, timing constraints, weather and traffic conditions, etc. So, normally you take a car to work, but if the weather is exceptionally bad the social practice does not force the default action, but rather gives input for deliberation about a new plan in this situation and take a bus or train (or even stay home).

In order to illustrate the major ideas of this paper we consider the development of a serious game to train first responders in a crisis management situation. The use case is a collision of a truck loaded with fluids with a car which has consequently caught fire. The use case is based on our experience designing a serious game for crisis management including such scenarios.

In Sect. 2 we describe how the prototypical approaches described above (goal directed, case based and work practice) would need to model the scenario and what are the main issues. In Sect. 3 we describe social practices and how they can be used in the deliberation cycle of agents. In Sect. 4 we describe how the scenario can be modelled, making use of social practices. Some conclusions and future work are given in Sect. 5.

2 Background

In this section, we discuss how activity in context is handled in different frameworks for social deliberation. In particular, we describe 2APL as a typical example of goal-directed approaches, Case-Based Reasoning as an exponent of situation based approaches and Work Practices as an approach in between. Of course, many other approaches exist that take (social) context into account when deciding about actions, but the approaches discussed here highlight the main issues.

2.1 2APL

The multi-agent programming language 2APL supports the implementation of individual agents that can perform high-level reasoning and deliberation about their information (i.e., beliefs) and objectives (i.e., goals to achieve) in order to decide what actions to perform [6]. In order to reach its goals, a 2APL agent adopts plans. 2APL provides programming constructs to implement beliefs, goals, actions, plans, events, and three different types of rules that can be applied to generate plans. In particular, 2APL provides *planning goal rules* that implement practical reasoning rules that can be used to generate plans for achieving

goals and *practical reasoning rules*, which can be used to expand abstract plans to concrete sequences of actions and to rewrite plans to cope with unforeseen circumstances [5].

A possible way to represent the top level planning goal rules for a fireman agent in the fire fighting scenario as indicated in the introduction is:

```
Goal: handlecrisis
GP/PR rules:
handlecrisis <- victim | inform(medics);save(victim);extinguishfire
handlecrisis <- | extinguishfire
extinguishfire <- chemical | spray(foam);clean
extinguishfire <- | spray(water);clean
```

The first rule states that whenever the precondition “victim” holds, the goal “*handlecrisis*” can be dealt with by the “*plan*”: *inform(medics);save(victim);extinguishfire*. The second rule states that *handlecrisis* can alternatively be handled by the plan *extinguishfire*. This specification makes use of the fact that in 2APL, rules are tried in order, and the first one that is applicable is executed. Thus the agent first checks whether a victim is present and in case there is, it will inform the medics, save the victim and extinguish the fire afterwards. If there is no victim, the condition of the first rule does not apply and the agent follows the second rule: extinguishing the fire. A similar process appears for the “*extinguishfire*” rules. It first checks for chemicals. If they are present it will use foam. If they are not present, it will try the second rule and use water.

A problem with this approach is that if the first rule fails during the execution of the plan, e.g. because the medics cannot be reached, the agent will try the next rule and start extinguishing the fire, even though there is a victim present. This can be avoided by explicitly indicating a precondition to be true or false in order to distinguish the different cases, as in the following two rules:

```
handlecrisis <- victim |
inform(medics);save(victim);extinguishfire
handlecrisis <- not(victim) | extinguishfire
```

However, if in this situation the plan associated with the first rule fails, the second rule is not applicable because there is a victim. In this situation, the agent would just stop, without saving the victim or extinguishing the fire, as it has no applicable plan to follow. Moreover, in the case of conditions involving several criteria, the number of rules would quickly increase such that the different combinations of conditions could be represented.

We are not claiming that the above cases could not in some way be represented in 2APL. However, the example highlights two aspects that are inter-related and mingled in the 2APL representation. The conditions of the rules function as a precondition of the plan in the rule. However, the same conditions are also used for rule selection. The latter necessitates the constructions shown above but can also lead to (unexpected) difficulties as indicated.

By using the idea of social practice we distinguish conditions that are needed for rule selection and preconditions of plans. Thus we do not incorporate the conditions in all the rules, but are checking the context conditions of the social practice first and given those conditions select a subset of the rules that are relevant for the situation. Thus the deliberation is no longer purely goal driven, but is goal plus context driven. This will lead to a more natural specification and (through the modularization of rules based on context) to a more efficient deliberation.

2.2 Case-Based Reasoning

Whereas 2APL follows a goal-based approach for selecting actions, Case-Based Reasoning (CBR) is an example of reactive deliberation. CBR uses previous cases (or situations) as the basis for the selection of the next action [17]. The general cycle of CBR follows the following steps:

1. problem formulation
2. retrieve
3. reuse
4. revise
5. retain.

The first step is to formulate the problem. This is important because the way a problem is formulated will determine the query on the case-base, through which the most relevant case is eventually selected. In our scenario this might lead to the following (simplified) problem formulation:

```
fire(house)= no
fire(tanker)= small
contents(tanker)= unknown
fire(car)= no
oil-spillage= no
victims(car)= 2
victims(tanker)= no
...
```

For simplicity, we use here a very simple attribute-value structure, but more complex structures can also be used. With this formulation the case-base is searched for a similar situation. The likelihood of finding a case exactly like the current situation is minimal. Thus one needs some metrics in order to find the most ‘similar’ case. Without getting into details, we just point to some difficult aspects here. Suppose there is an almost identical case except that it considers also a victim in the tanker. Could we use that case as a basis for the current course of action? As we indicated in the scenario it might be crucial to get information from the truck driver about the load of the tanker. If the driver is unconscious this is not possible, while if the truck driver is not injured he might give that information right away. So, even if the case in the case-base differs in only one parameter it might lead to a quite different course of action.

Given a case from the case-base that is close to the present situation, it has to be checked whether the plan for that situation can be used as it is or should be adjusted. E.g. if in the case from the case-base the first step would be to contact the owner of the tanker to ascertain the contents of it, this could now be just asked from the driver as he is not injured. Although for humans it is reasonably obvious how to make such a revision of the plan it is more difficult to find an algorithm that could calculate the necessary adjustment automatically. Finally, the system should decide whether the present case and its course of action are sufficiently different from the cases in the case base to warrant adding it (in the right place).

Intuitively the example makes clear that in scenarios like the one from Sect. 2 there are many parameters that potentially influence the course of action and even small differences can have big consequences for the course of action to be followed. Thus one needs to have a very large case base to cover all relevant cases such that an appropriate course of action is followed in each situation. In many domains (like crisis management) such a large case base cannot easily be assembled nor is it possible to construct one on the fly, because the consequences of errors are too big to allow for a gradually improving system.

2.3 Work Practice Simulation

A last approach relevant to social deliberation is that advocated by the Brahms platform [19]. Brahms is a multi-agent, rule-based, activity programming language. The Brahms language allows for the representation of situated activities of agents in a geographical model of the world. Situated activities are actions that happen in the context of a specific situation, thus their execution is constrained not only by the reasoning capabilities of an agent, but also by the agent's beliefs of the external world, such as where the agent is located, the state of the world at that location and elsewhere, located artefacts, activities of other agents, or communication with other agents and artefacts.

The philosophy of Brahms comes from the realization that work practices in organizations differ from the work flows as described and prescribed by the organization. If it is recognized that ultimately employee behaviours, rather than management practices, are the key to success in organizations [4], then these practices should be described as agent behaviours rather than the official (goal directed) plans. Within Brahms a work practice is defined as the (collaborative) performance of collective situated activities of a group of people who collaborate and communicate, while performing these activities synchronously or asynchronously, by making use of knowledge previously gained through experience in performing similar activities. Differences between formal plans and the work practice can lead to unforeseen results and render organizational plans useless.

The Brahms modelling language is geared towards modelling people's activity behaviour [19]. The Brahms framework consists of several interrelated models. Of particular relevance for this paper is the Activity Model that defines the behaviour of agents and objects by means of activities and workframes. Brahms has an activity-based subsumption architecture by which an agent's activities can be

decomposed into sub-activities. Activities can be interrupted and resumed, just as humans can multi-task by switching between different activities. Workframes control when activities are executed based on the beliefs of the agent, and on facts in the world. However, as in CBR, workframes require the full instantiation of all its preconditions in order to be applied.

Our scenario could be modelled through the use of several workframes. One for saving victims, one for checking the tanker load and one for extinguishing the fire. By giving the saving victims framework a high priority it will try to execute first. Thus if the preconditions of the workframe are fulfilled (which will include the presence of victims) it will start saving victims. If the workframe cannot be executed the one with the next highest priority will be executed, etc. If during the extinguishing of the fire suddenly a victim is discovered the framework for saving the victim is automatically fired and the extinguishing workframe is interrupted. This makes the Brahms framework quite flexible. However, it has the same problem as the 2APL framework in that the context and preconditions of the workframes are mixed. Moreover, the agents also do not have a learning capability that might lead to a priority adjustment of workframes after the (failed) execution of activities.

3 Social Practices

3.1 Social Intelligence and Social Practices

From the previous section, in which we showed some possible problems of using existing techniques and platforms for modelling social deliberation, as exemplified by a crisis management scenario, we can derive the following two requirements for socially realistic behaviour of systems active in real-time environments.

Context and preconditions for action. Socially intelligent agents should be able to understand and consider internal (pro-active) drives and external (reactive) drives. In most systems created for social interactions with humans, speed and appropriateness of reaction are leading for the system's behaviour. However, in open, dynamic situations, the socially intelligent agent also has to reason and plan for pro-active drives. So, the framework should maintain both internal and external drives and be able to reason about their relative importance in each situation. Whereas 2APL is strong on the goal based deliberation, it is difficult to incorporate the situation based reactions. This shows in the mix-up of context conditions that are situation based into the pre-conditions of plans for goals. CBR is strong in reasoning from situations, but lacks the goal based deliberation. Brahms seems to be better suited to model work practices, but it lacks flexibility in denoting situations and especially does not support learning.

Learning. The fact that intelligent agents should be able to learn is obvious. That most agent systems cannot really learn is also a fact. Considering socially intelligent agents, learning is of utmost importance given the need

for operation in highly dynamic and open environments and diverse interaction partners. The agent should use its experience to tune its parameters and learn the most effective actions in each situation taking into account both its physical and social effects. However, 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.

In order to cope with the complexity of combining situation based reactions and goal based planning while taking care of both social and physical aspects of reality and planning in such an environment, we propose to use *social practices* [16]. As described in the introduction, social practice theory seeks to determine the link between practice and context within social situations. That is, social practices aim to integrate the individual with his or her surrounding environment, assessing how that context relates to past, common, experiences, culture and capabilities of the individual. It should be emphasized that social practice theory is a sociological theory that takes social practices as focal point to explain social phenomena. Individuals only play a role in as far as they are ‘recruited’ by social practices in order for the social practice to be executed. In this paper we look at social practices from the individual’s perspective. It should be seen as an addition to the sociological theory not as an explanation or change of that theory. So, social practices certainly seem a good starting point for systems that need to take context into account.

Social practices can be seen as patterns which can be filled in by a multitude of single and often unique actions [16], that endure between and across specific moments of enactment. Through (joint) performance, the patterns provided by the practice are filled out and reproduced. Each time it is used, elements of the practice, including know-how, meanings and purposes, are reconfigured and adapted [18]. Therefore the use of social practices includes a constant learning of the individuals using the social practice in ever changing contexts. In this way social practices guide the learning process of agents in a natural way.

In [18] the social aspect of social practices is emphasized by giving the social practice center stage in interactions and letting individuals be supporters of the social practice. It shows that social practices are shared (social) concepts. The mere fact that they are shared and jointly created and maintained means that individuals playing a role in a social practice will expect certain behaviour and reactions of the other participants in the social practice. Thus it is this aspect that makes the social practices so suitable for use in individual planning in social situations. Because, in this paper, we concentrate on the individual planning we will not see much of the particular social aspects of the social practices, but we like to emphasize that it forms the basis of the success of the individual planning with social practices.

3.2 Characteristics of Social Practices

Researchers in social science have proposed a representation of social practices based on three broad categories [12]: materials, meanings and competences.

Based on these ideas, we developed a model to represent social practices that can be used in social deliberation by intelligent systems. Obviously, as is the case with e.g. the representation and use of norms, other representations of social practices are possible given the many dimensions of the use of social practices. Our proposal, depicted in Fig. 1, is especially suitable for use in agent reasoning.

The components of this representation model are as follows:

- *Physical Context* describes elements from the physical environment that can be sensed:
 - *Resources* are objects that play a role in the practice such as fire hose, fire, truck and car in the scenario.
 - *Places* indicates where all objects and actors are located relatively to each other, in space or time.
 - *Actors* are all people and autonomous systems involved, that have capability to reason and (inter)act.
- *Social Context* contains:
 - *Social Interpretation* determines the social context in which the practice is used.
 - *Roles* describe the competencies and expectations about a certain type of actors.
 - *Norms* describe the rules of (expected) behaviour within the practice.
- *Activities* indicate the normal activities that are expected within the practice. Not all activities need to be performed! They are meant as potential courses of action.
- *Plan Patterns* describe usual patterns of actions defined by the landmarks that are expected to occur.
- *Meaning* refers to the social meaning of the activities that are (or can be) performed in the practice. Thus they indicate social effects of actions
- *Competences* indicate the type of capabilities the agent should have to perform the activities within this practice.

Looking at the characteristics of social practices as given in Fig. 1 one can notice some resemblance to the aspects that also play a role in agent organization models (see e.g. [11]). This list can be seen as an analogue of the connection between imposed and emerging norms. Both organizations and social practices give a kind of structure to the interactions between agents. However, organizations provide an imposed (top-down) structure, while the social practices form a structure that arises from the bottom up. Thus where organizational interaction patterns indicate minimal patterns that agents should comply with, the patterns in a social practice indicate minimal patterns that can and are usually used by the agents.

In the next sections, we sketch how we envision the use of social practices in agent deliberations and how they indeed seem to be a useful part of a new architecture for socially intelligent agents.

Abstract Social Practice	Combat Fire	Going to work
Physical Context		
Resources	inflammable objects, water, barriers...	Vehicles, money, ...
Places	locations of fire, actors and resources,...	Stations, roads,...
Actors	Victims, bystanders, colleagues,...	Drivers, co-passengers, ...
Social Context		
Social interpretation	Dangerous places, safe places, rescue equipment	Bus driver, train driver...
Roles	police, medics, ...	
Norms	own safety; public safety	
Activities	Identify type of fire; Extinguish fire; Remove victims; Clear area; ensure own/team safety; ...	Choose transport type; buy ticket; drive car; ...
Plan patterns	<pre> graph LR Start[Start] --> FoundCasualty[Found-casualty] FoundCasualty --> SafeCasualty[Safe-casualty] SafeCasualty --> RescuedCasualty[Rescued-casualty] FoundCasualty --> FoundFire[Found-fire] FoundFire --> TeamFormed[Team-formed] TeamFormed --> ExtinguishedFire[Extinguished-fire] ExtinguishedFire --> End[End] </pre>	
Meaning	braveness, leadership, ...	Environmental conscientiousness, comfort, social status, ...
Competences	<ul style="list-style-type: none"> • Fire combat knowledge and skills • Coordination skills 	<ul style="list-style-type: none"> • driving skills, • cycling skills, • knowledge of public transport routes

Fig. 1. Social practices

3.3 Social Practices in Deliberation

Just by using social practices does not necessarily make agents become socially aware. We also have to have an agent deliberation in which social practices are taken into account at the right moments. In [8, 10] we sketched such an agent architecture. We will not describe it fully here again, but just highlight the aspects that are most important for the planning. We assume that sensing is not a passive but also an active process. Active social practices direct the agent's sensing to find objects, actors and events that are expected within that social practice. This leads to reactive but focused sensing based on the current situation.

When the interpretation of the sensed environment in terms of the existing social practices, results in only one possible action, the agent will perform that action directly. However, in cases that there are several possible courses of action, the agent will take into account its motives in order to determine possible goals that are applicable to the sensed environment and generate a plan accordingly. This deliberation can be a complex BDI deliberation or extensions thereof, such as the FAtiMA [7] or BRIDGE [9] deliberations containing emotions, goals, intentions, beliefs, roles, identity, etc. However, due to the context of the social practice the agent can limit the portion of the applicable rules and beliefs and also can use specialized rules for planning (patterns) that are (only) applicable within the social practice.

Alternatively, the drive to search for patterns (and thus the sensing process) can be steered from the agent's motives. E.g., a fireman with a high achievement

motive and low avoidance motive will start approaching a fire right away looking for a victim to establish whether a victim can still be saved. Thus we see that the parallel tracks of pro-active and reactive behaviour already start with the sensing behaviour.

Using social practices in deliberation, one can also distinguish the fast and slow reasoning tracks as described by Kahnemann [13]. If a social practice matches the features of a situation to an extent that it dictates a cause of action right away, this leads to a reactive action. E.g. seeing a small fire leads a fireman to use a fire-extinguisher and extinguish the fire (instead of starting to connect the water hoses).

A second fast track is taken when a social practice in combination with a motive also leads to a motivation for a certain type of behaviour. E.g. if a fireman has a strong achievement motive to have all victims saved and it gets the order to proceed from the fire commander he will get to the victim and start to evacuate her (even though he might not have planned the whole activity yet).

If the course of action is not completely determined by the social practice some more *deliberation* takes place.

Figure 2 gives a more detailed overview of the deliberation process.

As in traditional BDI reasoning, by perceiving the current context, an agent will revise its beliefs and goals. The context will also trigger some social practices. That is, some elements of a social practice are filled out by the sensed observations, resulting in more concrete social practices. For example, in the context of the crisis management scenario, the social practice ‘*rescuing-victim*’

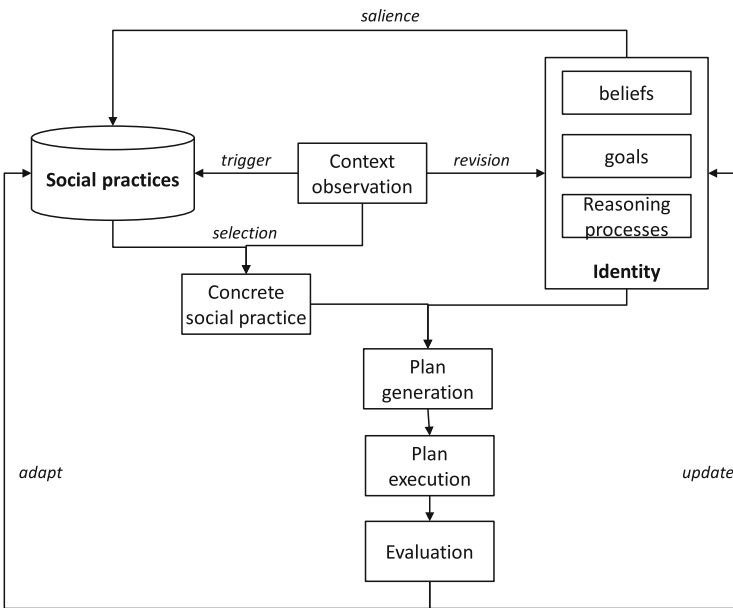


Fig. 2. Social reasoning process

is triggered. Further observation of the context will determine whether danger of explosion exists, whether other rescue workers are available, whether the victim is conscious, resulting in the more concrete social practice of '*rescuing-unconscious-victim-in-explosion-danger-together-with-colleague*'. If the agent can establish that the colleague has the same information or because they use a signal when getting in to synchronize the agent can assume that the colleague follows that same social practice and the rescue can be done efficiently without negotiating the coordination (because that is included in the social practice plan patterns). The agent will then generate a plan for this social practice based on its identity (beliefs, goals and reasoning process). Thus it should be emphasized that the social practices do not replace the traditional deliberation of BDI agents. They give a background and patterns that can be used for the plan deliberation. This will increase efficiency and also allows for dividing context checks and pre-conditions of actions and plans.

Notable in the above scenario is that the social practices take a leading role in organizing possible courses of action. Thus we do not have either a fixed set of plans per goal nor do we have a large set of plans that have to be searched. Because the social practices combine material and social aspects 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 also can instantiate elements in the deliberation if they are already clear from the context, such as the roles and expected goals. Other times social practices overrule elements like emotions thus bypassing deliberation about this element.

In order to facilitate their use, social practices should be stored in an efficient and easy to use way. Case Bases such as used in CBR provide seem to be suitable structure for the management and maintenance of social practices. The practices can be linked based on some (prominent) features and generalizations. Thus, for example, the relation between practices related to crisis management is based on the type of incident and parties involved. Generalizations also play a major structuring role. E.g. all fire fighting practices can be related to a general fire fighting practice. This allows for inheritance and all the usual reasoning over hierarchies.

Given the fact that agents will have many social practices, the question also arises how they choose a practice when several practices seem applicable in a situation. Several strategies can be designed for these situations. One is that the agent checks which social practice will most likely further its own (social) goals and motivations. The agent can also check its experience and choose the social practice that led most often to a successful interaction. So, it plays on safe. Many more strategies can be designed to make the choice. It is important to note that there is not a single possible best strategy to make this choice. The social practices function more as background and guideline than as obligations. Thus it is not necessary that in each situation a unique social practice fits. If more practices fit one will be chosen and its effects evaluated. If it worked well and the choice was based on an aspect not yet considered before, this aspect

might be added (conditionally). In this way the social practices form a flexible and evolving structure to support the adequate deliberation of the agents.

4 Scenario

In this section we illustrate how social practices can be used to model and implement the scenario of Sect. 2. When the fire brigade gets to the accident it might have heard about the situation already and assume it is in the *tanker fire* practice. If the firemen did not hear this yet, the first thing they will do is check which objects are on fire (this might be part of the social practice of “*arriving at an incident*”). The concrete social practice of “*tanker fire*” is exemplified in Fig. 3.

In this case, the firemen can start right away using this social practice of tanker fire to deliberate about their next actions. Notice that they might do this as well even if they have not seen the car that has been hit. As the situation unfolds they will continuously check whether the physical resources present match those expected give the social practice. The information about physical

Concrete Social Practice	Tanker Fire
Physical Context	
Resources	Tanker, car, fire truck, water resource,
Places	Geometric position of all objects
Actors	Truck driver, car driver, passenger, firemen, policemen, bystanders
Social Context	
Social interpretation	Victim in danger, explosion risk, chemical health hazard, safe distance
Roles	Victim, firemen, tanker driver, police, bystander
Norms	police takes care of road safety Fireman determines risk of explosion Fireman has power to order evacuation Fireman has power over police , victim and medics
Activities	Identify type of fire; Extinguish fire; Removal victims; clear area; ensure own/team safety; close roads; first aid to victim,...
Plan patterns	Determine truck load before evacuate victims Evacuate victims before extinguish fire
Meaning	Safety first
Competences	<ul style="list-style-type: none"> • Domain knowledge and skills: know type fire • Coordination skills : know who to call about fire/victims Choice/deliberation skills: <ul style="list-style-type: none"> • When explosion high choose to evacuate all • When victim quick to help choose free victim first • When fire extinguished choose clear debris • When situation safe choose open roads • ...

Fig. 3. Concrete social practice tanker fire

resources, places and actors can also give rise to (additional) sensing in order to fill in parameter values and objects that are present.

Given the physical context asserted, the social context is filled as far as possible. Some of the social context is given by the social practice. E.g. a person sitting in a car that has collided with the tanker that is now on fire will be a “victim”. However, the social practice also prescribes that an interpretation should be made that indicates whether the situation should be classified as “*chemical health hazard*”. Again this might be simple if the truck load is known, but otherwise might give rise to an explicit process to ascertain the truck load and whether this load can give rise to a chemical health hazard.

The social context also indicates a number of aspects that will be important for the execution and coordination of actions from the social practice. The roles indicate the type of actors that potentially interact within this practice. The norms indicate constraints and expectations of (inter)actions within this social practice. E.g. the firemen expect the police to ensure the safety of the road. This means that they will not check this safety whenever the police is around.

The next consideration in the social practice is the competencies that are needed to execute the actions within this social practice. E.g. in order to ascertain whether there is a chemical health hazard the fireman needs to have knowledge about chemicals and particularly about the dangers of these in high temperatures. If the social practice indicates that the first thing to do is to ascertain chemical health hazards and the fireman has too little knowledge about chemicals or no knowledge about the truck load he needs to take action to remedy this lack of knowledge. Also, if the fireman has no experience with extinguishing certain types of chemicals and has no complete knowledge he might decide to just evacuate the scene instead of trying to extinguish the fire. Thus in this way

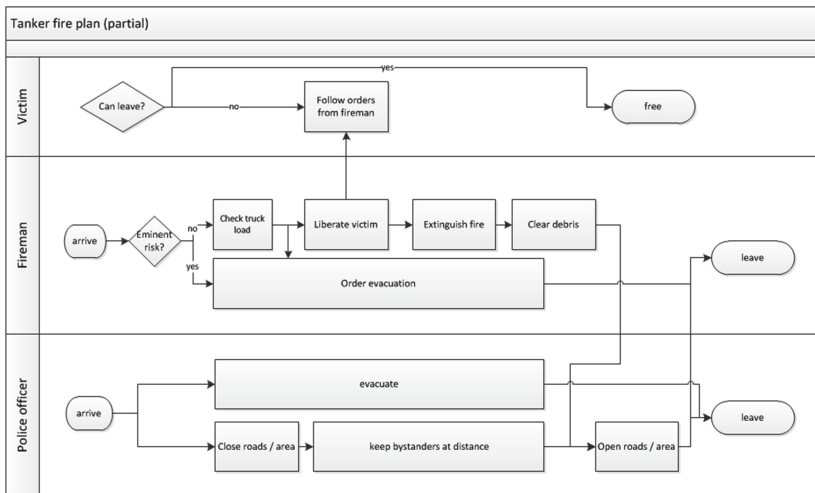


Fig. 4. Concrete plan

the competencies influence the particular actions that are executed or avoided within the social practice.

Finally, the plan patterns indicate that the fireman first determines the chemical hazard before evacuating victims. The victims have to be evacuated before he starts extinguishing the fire. These patterns determine a kind of default patterns that are assumed to hold for all ways in which the social practice is used. Note that the patterns in this case only give very sparse constraints on the behaviour. Thus the fireman has a lot of freedom to fill in the actual plan resulting from the social practice. Given his personality, experience, goals and emotions he can deliberate to come to a concrete plan as shown in Fig. 4.

5 Discussion and Conclusion

Intelligent agents increasingly need to be more socially aware of their context in order to take the appropriate action. We have shown how current techniques and platforms only provide partial solutions for this problem through the use of these approaches in a simple scenario. Of course, this does not mean that current platforms are not usable, but rather that they have problems modelling a scenario in which aspects are present that play a prominent role for socially intelligent behaviour.

In this paper, we have argued for the use of social practices as part of a new agent architecture that should facilitate socially intelligent behaviour. It puts social context and social motives at the heart of the deliberation rather than use them as additional modules. We have sketched how this architecture provides some structure in the complexity of the deliberation process, facilitates the combination of social and physical aspects of a situation, integrates fast and slow thinking patterns as described in the psychology literature, and balances between pro-active and reactive behaviour. As such it can be seen as combining the features of goal directed (BDI) architectures as exemplified by 2APL, situation based reasoning as performed in CBR and workframe based deliberation as done in Brahms.

The use of social practices for social intelligent agents also led to the realization that the concept is not very well defined and thus needs to be made more precise and formal in order to serve as a basis for implementations. In this paper we have shown a first step towards this goal. Two aspects are important in this respect. First, social practices can be seen as a kind of emerging organizations. Taking this seriously led to the realization that many concepts of organization models can be very well used to describe aspects of social practices as well. However, one should bear in mind that social practices do not have the kind of imposed normative flavour but rather have the flavour of emerging norms.

The second aspect that should be further explored is the way social practices should be structured with respect to each other. Do we need abstraction hierarchies? As a first step it seems that using similar structures as used in case based reasoning is promising. However, more research with large quantities of practices should be done to make this more precise.

Because in this paper we have focussed on the use of social practices by the agents we have not touched upon the fact that social practices emerge and evolve in social interactions. Therefore an important characteristic of them is that they are partially shared. This means that when a fireman starts the tanker fire practice it can expect a certain type of behaviour from the policeman involved as they share the same practice. A shared social practice thus can serve as a background for facilitating coordination and also for solving problems in coordination protocols (due to a changed environment or social context).

References

1. Barnett, E., Casper, M.: A definition of “social environment”. *Am. J. Public Health* **91**(3), 465–465 (2001)
2. Boissier, O., van Riemsdijk, M.B.: Organisational reasoning agents. In: Ossowski, S. (ed.) *Agreement Technologies. Law, Governance and Technology Series*, vol. 8, pp. 309–320. Springer, Dordrecht (2013)
3. Bordini, R.H., Hubner, J.F., Wooldridge, M.: *Programming Multi-agent Systems in AgentSpeak using Jason*. Wiley, Chichester (2007)
4. Colvin, A., Boswell, W.: The problem of action and interest alignment: beyond job requirements and incentive compensation. *Hum. Resour. Manag. Rev.* **17**, 38–51 (2007)
5. Dastani, M.: Modular rule-based programming in 2APL. In: Giurca, A., Gasevic, D., Taveter, K. (eds.) *Handbook of Research on Emerging Rule-Based Languages and Technologies: Open Solutions and Approaches*, pp. 25–49. IGI Global, Hershey (2009)
6. Dastani, M.: 2APL: a practical agent programming language. *Auton. Agents Multi-agent Syst.* **16**(3), 214–248 (2008)
7. Dias, J., Mascarenhas, S., Paiva, A.: FATiMA modular: towards an agent architecture with a generic appraisal framework. In: Bosse, T., Broekens, J., Dias, J., van der Zwaan, J. (eds.) *Emotion Modeling. LNCS*, vol. 8750, pp. 43–55. Springer, Heidelberg (2014)
8. Dignum, F., Dignum, V., Jonker, C., Prada, R.: Situational deliberation; getting to social intelligence. In: *Computational Social Science and Social Computer Science: Two Sides of the Same Coin*, June 2014
9. Dignum, F., Dignum, V., Jonker, C.M.: Towards agents for policy making. In: David, N., Sichman, J.S. (eds.) *MAPS 2008. LNCS*, vol. 5269, pp. 141–153. Springer, Heidelberg (2009)
10. Dignum, F., Prada, R., Hofstede, G.J.: From autistic to social agents. In: *AAMAS 2014*, pp. 1161–1164, May 2014
11. Dignum, V.: A model for organizational interaction: based on agents, founded in logic. Ph.D. thesis, SIKS Dissertation Series 2004–1, Utrecht University (2004)
12. Holtz, G.: Generating social practices. *JASSS* **17**(1), 17 (2014)
13. Kahneman, D.: *Thinking, fast and slow*. Farrar, Straus & Giroux (2011)
14. Minsky, M.: A framework for representing knowledge. In: Collins, A., Smith, E. (eds.) *Readings in Cognitive Science*, pp. 156–189. Morgan Kaufmann, San Mateo (1988)
15. Panagiotidi, S., Vázquez-Salceda, J., Dignum, F.: Reasoning over norm compliance via planning. In: Aldewereld, H., Sichman, J.S. (eds.) *COIN 2012. LNCS*, vol. 7756, pp. 35–52. Springer, Heidelberg (2013)

16. Reckwitz, A.: Toward a theory of social practices. *Eur. J. Soc. Theor.* **5**(2), 243–263 (2002)
17. Richter, M.M., Weber, R.O.: *Case-Based Reasoning: A Textbook*. Springer, Berlin (2013)
18. Shove, E., Pantzar, M., Watson, M.: *The Dynamics of Social Practice*. Sage, London (2012)
19. Sierhuis, M., Clancey, W.J., van Hoof, R.J.J.: Brahms an agent-oriented language for work practice simulation and multi-agent systems development. In: Seghrouchni, A.E.F., Dix, J., Dastani, M., Bordini, R.H. (eds.) *Multi-Agent Programming*, pp. 73–117. Springer, New York (2009)
20. Smolka, A.L.B.: Social practice and social change: activity theory in perspective. *Hum. Dev.* **44**(6), 362–367 (2001)
21. VanHee, L., Aldewereld, H., Dignum, F.: Implementing norms? In: Hubner, J.F., Petit, J.-M., Suzuki, E. (eds.) *Proceedings International Joint Conference on Web Intelligence and Intelligent Agent Technology*, pp. 13–16 (2011)