

UTRECHT UNIVERSITY

Bachelor's thesis Artificial Intelligence (7.5 ECTS)

Implementing Social Practices for Agent Systems

Author:

R.J.J. BROUWER

Supervisor:

Dr. F.P.M. DIGNUM

June 26, 2015

Abstract

In this report I discuss the sociological theory of social practices and its application to agent systems. The central question that this report attempts to answer is whether it is possible to provide a concrete and precise description of the vague concepts related to social practices. The central concepts of the theory are discussed at length and some issues concerning their implementation are explored, based on a representational model proposed by Dignum and Dignum (2014). These ideas are further elaborated upon, leading to an investigation of two particular processes: *recognition* (or activation) of a social practice and *deliberation* based on social practices. Finally, a partial architecture for agents systems using social practices is proposed, including some limiting conditions for the remainder of the architecture.

Table of Contents

1	Introduction	1
2	Scenario	2
3	Social Practices	2
3.1	Social theory	2
3.2	Representation model	4
3.3	Towards an implementation	13
3.4	Modified model	15
4	Applying concepts	15
4.1	Working out the scenario	16
4.2	Prerequisites	16
4.3	Recognition	17
4.4	Plan composition and action selection	20
5	Architecture	24
5.1	Abstract architecture	24
5.2	Loose ends	27
6	Conclusions	28

1 Introduction

The decisions we make every day are heavily influenced by the social context in which we make them. Would you run a red light late at night if there was no-one around, except for one man standing at the corner? What if the man wore a police uniform?

Humans generally take the social context — the physical and social setting in which something takes place — into account when they plan goals or actions. Artificial systems lack such advanced social intelligence: traditional approaches are often unable to support truly social behaviour [2, 5]. This calls for a new architecture for agent systems that places social principles at the heart of the system [5, 6].

I will discuss several aspects of such an architecture that uses social practices¹ to describe social context. Social practice theory is a sociological theory that attempts to explain action and social order. A comprehensive discussion of this theory can be found in [12].

Social practices could serve as the basis for the deliberation process of an agent, the process where an agent decides on its course of action. Using social practices as a basis ensures that this decision process is inherently social.

In this report, there will be no discussion about the relative merits of this approach for modelling social behaviour. This has been done in [5, 6], where the first steps towards an architecture integrating social practices have been made as well. These papers will serve as the basis of this report. I will discuss the concepts associated with social practices at length, and will consider some aspects concerning the implementation of social practices. I will not provide a complete architecture for this purpose, but I will take steps towards such an architecture, and provide a basis for future extension and elaboration.

This report will thus take the first steps towards an implementation of social practices in agent systems and will attempt to provide an answer to the following question: is it possible to provide a concrete and precise description of the vague concepts related to social practices?

First, a scenario will be introduced in section 2 that will serve as illustration in the following sections. The key concepts of Social Practice Theory will be discussed in section 3, and I will work towards a model that can be used in an abstract architecture. Following that, section 4 will elaborate on the associated processes, and work towards a basic architecture presented in section 5. Finally, I will draw some conclusions and make suggestions for future research in section 6.

¹According to [12], the term ‘social practice’ is a tautology. I will use this term nonetheless, to avoid any confusion with other meanings of the word ‘practice’.

2 Scenario

Imagine a manifestation of unhappy employees on strike at the gates of a car factory. They are protesting an upcoming salary reduction by their employer. As the risk of a confrontation with employees that continue their work at the factory is considered substantial, a certain police presence is maintained.

Initially, no problems arise. The strikers confront the employees entering the factory, but let them through after a brief discussion. At some point, however, a discussion between one of the workers and the strikers gets out of hand. A scuffle arises and the police intervene. At this point, the strikers collectively turn against the police. This results in a small fight between the protesters and the police. Several arrests are made and eventually the crowd disperses, ending the manifestation.

3 Social Practices

Social Practice Theory has been proposed as the basis for a model of social intelligence in multi-agent systems [6]. In this section, I will discuss several aspects of Social Practice Theory, starting from the sociological theory. I will proceed with the discussion of a more precise representation model for implementing social practices in multi-agent systems, discussing each of the concepts in detail along the way.

3.1 Social theory

Social Practice Theory is, originally, a sociological theory, which means that it concerns large groups of people and their (inter)actions. Individuals are merely carriers of a social practice or agents that perform social practices. The individual is seen as the intersection of the practices it carries [12].

This approach is not suited, nor meant, for the purpose of describing the behaviour of individual agents. A shift from the group perspective to the individual's perspective is necessary, if the theory is to be used as the basis for an agent architecture. Agent systems are distributed autonomous systems, which entails that we can hardly see social practices as universal concepts that 'recruit' agents. Agents need to have their own representation to use in deliberation and plan generation, independent of other agents or an external superimposed structure.

As individual agents (ideally) have the ability to learn, there is no guarantee that social practices are identical in different agents. This is exactly what is needed, as the idealized universal social practice is both unrealistic and impractical: there generally is variation among a population in the way social situations are handled.

If this were not the case, the (social) individual would become a rather empty notion. It is thus expected that the representation of a social practice may vary somewhat among agents, based on individual experience. With this in mind, I will discuss the concepts of Social Practice Theory next. The discussion is based on the concepts as described by Reckwitz in [12].

Reckwitz defines social practices as follows:

A ‘practice’ (*Praktik*) is a routinized type of behaviour which consists of several elements, interconnected to one other: forms of bodily activities, forms of mental activities, ‘things’ and their use, a background knowledge in the form of understanding, know-how, states of emotion and motivational knowledge. A practice (...) forms so to speak a ‘block’ whose existence necessarily depends on the existence and specific interconnectedness of these elements, and which cannot be reduced to any one of these single elements. [12]

These elements can be classified in three broad categories, on which social practices are based: *materials*, *meanings* and *competences* [9]:

- *Material* concerns all physical aspects of the social practice. Of the elements Reckwitz names, ‘things’ and *forms of bodily activities* probably fall in this category, and possibly *forms of mental activities* as well. Some examples taken from the scenario described in the previous section would be: a megaphone (a ‘thing’, used by the employees to voice their discontent), a policeman (an agent) and the halting of an employee entering the factory (an activity).
- *Meaning* concerns everything that is connected to the physical level in terms of interpretation, understanding, etc. Reckwitz’s, *states of emotion*, *motivational knowledge* and *the use of ‘things’* would fit here. Again, some examples from the scenario: the function of the megaphone (its use in this context), the conviction (of the employees) that their salary should not be reduced (a belief) and the anger they feel because of that reduction (an emotion).
- *Competence* concerns the skills and knowledge that are needed to perform actions within the social practice. From Reckwitz’s list, *background knowledge in the form of understanding* and *know-how* count as elements of competence. Examples from the scenario could be: the ability to converse with other employees (a skill) and knowledge of how to turn on a megaphone (knowledge).

The individual is the link between these concepts. Meaning and competence are inherent to the individual (although they are by no means fixed, but subject to learning). This individual, as a container for meaning and competence,

adopts (uses, perceives, executes) material, thus providing the link with the final component.

Although this conceptual framework already provides a structured representation of social practices, it remains very abstract and vague. This is easily realized when thinking about the placement of Reckwitz’s concepts in these categories. Some of them do not fit perfectly in any of the three categories, or might be placed in multiple categories. This classification is not suitable for describing (inter)action on an individual level. The framework may serve as a basis, but it needs to be elaborated and refined further. Typically, this is most effectively achieved when one starts to think about (a concrete) implementation of the theory. In the next sections, I will describe and discuss the first steps towards applying this sociological theory in computational practice.

3.2 Representation model

Earlier attempts at implementing social practices in multi-agent systems generally do not take a(n) (sufficiently) individual approach. Models for social practices have been proposed by Hoffman and Holtz [7, 9].

Hofmann came up with a simulation model, based on Social Practice Theory, to model and investigate repeated patterns of social behaviour [7]. He treats social practices as joint activities, and does not take the individual level into account at all.²

Holtz has proposed a model to investigate the generation of social practices [8, 9]. This model only implements social practices on a very basic and abstract level. It is not concerned with the involvement of social practices in the *deliberation* of individual agents, and thus does not suit our purpose either, although we have already seen the use of Holtz’s work in clarifying the sociological theory in the previous section.

Recently, a proposal for a model of social practices from the individual’s perspective has been put forward by Dignum and Dignum [5], building on earlier work sketching an agent architecture for this purpose [2, 6]. The discussion below will center around their representation model.

On the individual level, one could think of a social practice as a complex condition-action rule [6]. The first step is to match a situation to relevant social practices. In the scenario, one could see how a social practice such as *small-scale-protest* might be relevant to a policeman in this situation, while a social practice such as *large-scale-forest-fire* has no applicability at all. The main challenge for an

²“Individuelle Handlungen werden nicht berücksichtigt.” (p. 69); “Da das individuelle Verhalten der Agenten in dieser Simulationsstudie nicht berücksichtigt wird, sondern es sich um eine Simulation des Prozesses handelt, der dazu führt, dass sich in einer Gruppe von Agenten die geteilte Absicht herausbildet, eine gemeinsame Handlung durchzuführen, (...)” (p. 81).

Concepts	<i>Small scale protest</i>
Physical Context Resources Places Actors	baton, cuffs, signs, ... location of protest, location of police, ... colleagues, protesters, bystanders, ...
Social Context Social interpretation Roles Norms	enforcement equipment, protest zone, ... police, protesters, ... public order, freedom of speech
Activities	stand down, arrest aggressive protesters, call for backup, clear area, ...
Plan patterns	<pre> graph LR Start[Start] --> protest[protest] protest -.-> arrest[arrest] arrest --> clear_area[clear area] clear_area --> End[End] protest -.-> protest_end[protest end] protest_end --> End </pre>
Meaning	aggressiveness, service, defiance of authority, ...
Competences	identification of aggressive individuals; knowledge of local laws; defensive skills; ...

Figure 1: Schematic representation of the representation model by Dignum and Dignum [5].

implementation is to provide a quick but precise matching of situations and social practices.

During the selection of a relevant social practice, the material elements of the situation will get a social interpretation that is at least partly determined by the social practice (e.g. Reckwitz’s ‘use of things’ and roles of agents). The group of employees in front of the factory might be given social meaning as *protesters*, for example.

A social practice should also define several behaviours that are possible or desirable within it. The agent will select an action, taking into account its own competences/capabilities, and this action in turn has an expected social effect. A police officer in the *small-scale-protest* practice could, for example, *stand-down* or *arrest-aggressive-protesters*. The choice is influenced by detailed information about the environment and the competences of the police officer (the skill of choosing when either action is appropriate can be seen as a skill in itself). Initially, there is no reason to act. But as soon as the first fight starts, the presence of aggressive individuals triggers the police to choose the latter action.

Dignum and Dignum describe a representation model that captures all of these elements [5]. Fig. 1 provides a schematic representation of the model they propose,³ along with examples of each concept from the *small-scale protest* social practice. We will discuss all concepts in detail below.

Physical Context

The *physical context* in this model describes the elements in the physical environment. It represents part of the *material* category from the social theories, as it contains the objects, actors and locations from the physical environment (‘things’). It does not, however, completely cover this category, as the activities are not contained in it. This provides a better separation of the situation and performance of the social practice.

The three subgroups of the physical context each represent different aspects of the physical environment, that should be treated differently:

- *Resources* are objects that are relevant for the social practice. Cuffs can be seen as a resource in the *small-scale-protest* practice that I described earlier. A pillow would generally not be considered relevant to this social practice, and should not count as a resource, even if it is (always) available in the environment. In short, only objects that (potentially) play a role in the social practice should be mentioned here.
- *Places* are locations (in space and/or time) where the resources and actors are located. The spatial location that is occupied by the patrolling police officers is a location of significance in this social practice: it is risky for protesters to come near it, and they will usually keep some distance.
- *Actors* are all agents involved in the social practice. This includes autonomous systems and people. The key difference with resources is that agents are more than usable objects: they have the capability to reason and act by themselves. Note that, even though I already name the groups according to their role in fig. 1, role-assignment or identification by role is not part of this component.

Social Context

The *social context* places a ‘social layer’ over the physical environment. It provides an interpretation for objects, agents and locations. In the model under discussion, this is divided in three subcategories:

³Based on fig. 1 in [5].

- *Social Interpretation* provides that layer for everything - except for the actors in the practice. It provides social context for objects and locations but also for situations (e.g. “Employee threatened”).
- *Roles* describe the competences of certain actors, and what type of behaviour might be expected from them. It can be expected of a police officer to have better defensive skills than a civilian, for example. And one would sooner expect to be arrested by a police officer than by a civilian.
- *Norms* are rules that define (expected) behaviour within the social practice. An example of a norm would be that only the police is allowed to apply (controlled) violence.

Roles and norms are complex concepts in themselves. A role should contain at least (part of) the objectives of an agent enacting that role, its rights/capabilities and a set of norms/rules that define the (expected) behaviour of the agent. These three elements capture the goals of agent, the capabilities it has to reach these goals and the rules within which it should usually operate. Roles have been studied as concepts in an agent society before (e.g. [3]), and this might serve as an inspiration for the representation of roles for social practices.

It is not easy to determine at what level roles should be represented. In the running example, I have sketched a rather high-level view of the situation. The social practice I described, *small-scale-protest*, remains fairly abstract. The mentioned roles are no different: the role of *protester* would fit a fair share of the agents involved. Only a small fraction of those would, for example, fit the role of *union leader*. A *union leader* would have slightly different objectives (and possibly different capabilities and norms as well) compared to the other protesters. It is not a solution to exclude *union leaders* from the group of *protesters*, as they are definitely part of that group. They share certain objectives and norms, and other actors (such as the police) might not treat them any different.

A possible solution is to use inheritance in role modelling, and let *union leader* be an extension of *protester*. Such a solution was proposed for social contexts in [1]. The way roles relate to each other is very similar to the way social practices do and I will not discuss this extensively here. The notes about this in section 5 can be applied to roles as well.

The term ‘norm’ has a wide variety of interpretations. Norms in the context of this model are rules of (expected) behaviour. That is, rules to which agents generally comply. An example of a norm in the *small-scale protest* social practice is that one does not assault another (and does not expect others to do so). When the context shifts to a small riot after the police intervention, this norm is no longer applicable: a shift towards another social practice has taken place. Norms,

in this interpretation, can clearly vary across contexts, and thus require explicit representation in the social practice.

Additionally, norms may be role-specific: a bystander expects the police to arrest individuals that display aggressive behaviour, but he does not expect protesters to act in that way. Norms may be represented within the definition of a role (as is done in [3] for example), because norms are exclusively coupled with agents and not directly relevant for object or locations, as they do not execute actions or display behaviour. Defining norm at a role level has some advantages, but exclusively defining norms at this level leads to two problems. First, there is no way to represent the variation of norms for a role between contexts and, second, shared norms are overrepresented. Additional representation at the level of the social practice is necessary to capture variation of norms between contexts, and define shared norms once for all roles (“All agents should...”).

Norms can be represented in deontic logic. See for example [11] for a discussion of the use of deontic logic in normative systems.

The division of social interpretation for actors on the one hand and everything else on the other, as proposed by the model under discussion, leaves room for discussion. Roles are a fairly direct social interpretation of actors, while in *social interpretation*, the interpretation of objects, locations and situations is grouped together. *Social interpretation* therefore remains a rather vague concept. It is true that roles are probably the most complicated concept of all, but it might nonetheless be desirable to split this social interpretation, to provide a social counterpart for all concepts defined in the physical context. This would allow for clearer linking between the social and the physical.

Furthermore, it might not be needed to treat situations separately from objects and locations. Situations are bound to elements of the physical environment. They can be seen as parts of the social context of these elements. For example, “threatened employee” can be seen as a fairly specific role, or part of a more general role, “health hazard” as part of the interpretation of the objects or agents that pose a threat and “safe distance” as the interpretation of a location at a certain distance from a threat.

One might wonder why there is a distinction between the physical and the social within the social practice at all. Is it not the case that an actor is, within the practice, nothing more than its role? And is a resource not equivalent to its use within the practice, as a place to its interpretation? The actual physical agents, objects and locations are provided by the environment. The social practice adds social meaning to those objects. Why define physical context within the definition of a social practice at all?

The answer should be that the physical context within the social practice provides the link between the social interpretation and the representation of the environment. Take agents, actors and roles for example. A role is unique. There can only be one *police officer* role, yet it is possible that multiple police officers play a role in a social practice. When the physical context is removed from the practice, there is no way to represent this. It would be wrong to just add *police officer* as a role and be done with it - even if all police officers would act exactly the same. There can be a significant difference between practices that, on the social level, only differ from each other in the number actors of each role. An example taken from my own experience would be the difference between a lecture attended by ten students and one attended by three hundred. Norms differ between these contexts, and both students and teacher behave very differently. This difference cannot be captured by roles and social interpretation alone, as there is no difference between the two social practices on this level. Even though the social context is not completely identical, the norms are different, this is not enough to differentiate between the social practices. It is impractical to select social practices based on active norms, as this is not something an agent can (easily) perceive.

The physical layer provides the translation of agents and objects in the environment to roles and uses. But it does not just deal with the numbers and combinations of roles that occur in the practice, but also with what types of agents/objects can count as actors filling a particular role or resources having a particular use. Without the physical layer, we have agents/objects on the one side and roles/uses at the other, but no way to link the two. Take enforcement equipment for example. What counts as such? Assuming the agent recognizes these objects: batons and cuffs. The physical layer should provide this translation. It is clear that this is strongly dependent on the model of the environment that the agent maintains. If it does not recognize batons as such, the description must be given in terms that the agents *can* (directly) perceive: *black-stick-with-handgrip*, for example.

To summarize the above discussion: the physical context links social concepts to objects in the environment. It serves both as a dictionary (what can count as an actor with a certain role) and a many-to-many relation (how many agents of a role are expected and in what combinations). The physical context will form the basis of the selection of relevant social practices, which will be further discussed in section 5.

Activities

Activities are potential courses of action. It includes all activities that can be performed by the agent itself that can be expected within the social practice. Not all activities represented in a social practice *need* to be performed. They are possible courses of action, from which the agent can choose. The activities

contained in the practice will vary between agents: a protester most likely does not have *arrest-aggressive-protesters* available as a potential course of action, while a police officer does.

The activities available to an agent could be defined purely at the individual level – but it is clear that there is a strong link with the role of an agent. It might be beneficial to link activities with roles and let the agent think of itself as performing a particular role. This would also bring the agent’s concept of itself and of other agents (roles) closer together. This approach would reduce an agent to a performer of a role, which might be a rather constraining conception of an autonomous agent. This is partially remedied by the fact that the exact interpretation of a role may differ between agents, as the representation is individual, which allows for variation in role performance. Also, differences in competence (see below) may cause some variation in the choices made.

Another problem with this approach, that is less easily dismissed, is that agents cannot reason about other agents with the same role in another than their own interpretation of that role (i.e. they cannot differentiate between themselves and others performing a role). This might not be problematic in practice, but it limits the freedom of the agent to act outside of standard patterns. An example where this would be problematic is the following: a corrupt government official has a certain expectation of government officials in general. It is unlikely that he considers all government officials to be corrupt, rather, he flourishes because this is not the case. His self-conception is rather different from his conception of other government officials, but it would probably not be justified to consider the corrupt government official as having a different (personal) role than the others.

It is probably best to view activities only in relation with plan patterns. Individual activities may be performed only in the context of a plan. Activities are, as I will explain in the next section, part of the plan patterns. The plan patterns, in turn, have a strong link with the role performing them and, as we will see in section 4 form the basis for the selection of actions.

Activities might thus very well contain activities that the agent itself cannot perform, but are contained in a plan that is meant for a different role. The agent needs representations of such activities as well, as an understanding of these activities might be essential in the interaction with other agents.

Plan patterns

Plan patterns describe plans that the agent could follow. They structure sequences of actions, organized in scenes. Plan patterns can be of great help in plan generation, as they limit the number of possible activities that have to be considered. Additionally, plan patterns reduce the need for constant negotiation. When two police officers arrest an individual, they will presumably carry out the same plan

within the same practice. A police officer knows what has to be done to make a lawful arrest, and he also knows (or: expects) that his colleague probably has the same knowledge. This enables them to carry out the plan together without need for further negotiation, once the plan has been initiated. Plan patterns can then be seen as knowledge in the sense of *knowing what to do*. It is thus a competence, at least in the sense of [9].

It is important to notice that plan patterns are not (necessarily) meant as scripts. They do not define a fixed sequence of actions that will be executed as such once certain conditions are met. Rather they define (at least) the relative order of actions, e.g. a police officer should read someone his rights before asking him to give a statement about the behaviour that led to his arrest. It is not necessarily the case that no other actions may be executed in between. It is perfectly acceptable for a police officer to call for backup in between the actions mentioned earlier.

The representation of plan patterns by Dignum and Dignum is very similar to the *interaction structures* from section 3.4 of [3].⁴ This seems like a sensible choice: interaction structures are partial orderings of *scene scripts*, and scene scripts, in turn, describe the interaction of roles in a context in terms of activities. These scripts can be strictly or loosely specified: they can be used as protocols that the agent must follow, or just define some constraints and objectives [3]. This fits with the idea of plan patterns defining a partial ordering and constraints on actions.

The example in fig. 1 follows the conventions as laid out in section 3.4 of [3]. The blocks represent scenes. The scenes are not expanded in this example, but keep in mind that each block represents a scene script as described above. The dashed arrows in the example indicate exclusive choice, the solid ones represent inclusive choice.⁵

Meaning

Meaning is a concept that is strongly connected to plans and activities. It defines their social meaning and the effects of activities on the *social landscape*⁶. When the police would use (extensive) force in an intervention at the early stages of the protest, the protesters will probably consider this a rather aggressive action. Because of this, the respect they have for the police might decline, leading to a lowered acceptance of their authority, and a larger tendency towards defiance of

⁴A summarized discussion can be found in section 3.2 of [4].

⁵For the complete description of the syntax of the transitions, refer to section 3.4.2.2. of [3].

⁶Dignum et al. describe the social landscape as being the social counterpart of the physical landscape, indicating the social value of situations and agents. The spatial dimensions of this landscape denote social concepts such as desirability. Closeness in this landscape then indicates (social) similarity, which is a possible representation of social closeness [2].

that authority. All this is contained in the social meaning of that action. It should be clear that this is a rather complex and vague concept.

As meaning has a strong link with activities, it might be a logical step to link the two, i.e. make meaning part of the representation of activities within the social practice. This introduces some limitations as, intuitively, meaning is not just reserved for single actions. The meaning of an action may change due to the precise context and the actions that precede or follow it. The above example illustrates just this: if the police intervene early on, a neutral bystander might perceive this as an act of aggression, but when it is done at a later stage, at the start of a fight between protesters and employees, the meaning could change to one of sensibility and braveness. One could argue that a sufficiently different context warrants a different social practice, such that the meaning of an action within one social practice remains the same regardless of (small) context shifts. Large context shifts then lead to changes in the applicability of social practices.

This does not solve the problem entirely. As the order of actions might influence their meaning as well. Recall the situation where a police officer made an arrest. When he initially forgets to read an arrested individual her rights, and does this only after he has taken her statement, his colleagues will probably consider him a fool, and will not as easily trust him with the same responsibility in the future. This is an example of an agent not adhering to the plan patterns, but other types of examples are easily thought of.⁷

It is important that agents are able to reason about meaning. The aim of introducing social practices into agent architectures is to make agents aware of their social situation, and let them exhibit social behaviour. It is therefore essential that they are able to consider what the social consequences of an action are. Meaning should play a central role in the process of selecting actions.

Competences

Competences are the capabilities an agent should possess to be able to perform the activities within the social practice. It closely resembles the concept of *competence* from section 3.1. Competences include all types of capabilities: skills and know-how, physical and mental abilities. Competences (for policemen) in the social practice under discussion might be the knowledge of the procedure of arresting someone or the skill/ability to ward off attacks of aggressive protesters.

⁷Of course, whether such examples contain a ‘violation’ of a plan pattern as well depends entirely on the exact representation of the plan patterns. If these patterns are very precise (almost like scripts) such that they allow effectively only a very small set of action sequences, most examples will be of the same type. If this is not the case, a different example might be extending your hand to someone at your first encounter. This is generally considered a polite and friendly gesture. But if you wipe your nose with that hand shortly before extending it, the meaning changes significantly.

Some considerations need to be taken into account when defining capabilities. It is a rather complicated issue to determine when an agent has *sufficient* skill to perform a certain action. At what point does a policeman have sufficient skill to ward off attacks? It is hard to define this. The ambition should be to define competences as precise as possible (e.g. ‘has completed self-defence course’ instead of ‘can ward off attacks’) and they should be considered *minimal requirements* instead of *desired qualities*. This entails that it is much more straight forward to say whether an agent has the required capabilities or not. If it does not meet all requirements, it simply does not have the required competences.

Competences are very closely linked to activities. Every activity could be extended with ‘minimal requirements’ or preconditions (at the competence level). This is not problematic here as it was with meaning, as the required competences for an action hardly ever change.⁸

3.3 Towards an implementation

I will now summarize the most important issues from the discussion in this section. The main points that will play a key role in the design of a model and architecture for social practices in the next sections are:

- The physical layer provides the translation from perceived objects and agents to the (social) concepts of the social practice: both in number and type;
- Concepts in the physical layer should mirror their social counterparts;
- Norms should be represented both at the level of roles and at the level of social practices;
- Plan patterns provide constraints on action sequences;
- Meaning can be associated with both individual activities and sequences of actions (or plan patterns);
- Competences are closely linked to activities.

There is also an issue concerning social practices as a whole that needs to be considered: what level of detail do social practices represent? When is something considered a separate social practice? And when is it just another possible plan within a more generic social practice? This question is not easily answered: even on a case-to-case basis it proves hard to provide an answer.

⁸Except, maybe, when taking co-operation in account. An agent that is unable to communicate with someone in the English language might for example be able to do so when co-operating with another agent acting as a translator.

Take for example the social practice of dining out. Consider the difference between dining with a vegetarian and dining with someone without any dietary preferences. If the vegetarian is very strict (i.e. he does not want to eat food that has been prepared in a kitchen where meat is processed as well), it will alter the plan significantly and it might be justified to create a separate practice. This might solve part of the problem, but it creates a new one. Unfortunately, the real world is not black-and-white. It is rather difficult to decide when either of the practices apply. Suppose I will dine with a vegetarian that has no problem with eating in regular restaurants. What social practice applies? The one we created for the strict vegetarian does not seem to apply, so the general practice applies. But at what point should the switch to the other practice be made? This is not clear.

It will always be difficult to decide the applicability of practices, as it is nearly always impossible to delimit the regions where practices apply precisely. The creation of separate practices for increasingly specific situations increases the severity of this problem. At design time, one should make sure that the social practices are mutually exclusive.⁹

The level at which the social practices should be defined is strongly dependent on the application for which the agent is designed. An agent with a very specific purpose will probably benefit from very specific social practices within its ‘area of expertise’, while an agent with a broader purpose cannot handle that level of detail for practices, as the number of practices would be enormous if it is to function in many contexts.

It seems that the level of detail that social practices should represent has a strong link with the level of detail that the agent should be able to handle. Although this seems obvious, it requires careful consideration when defining social practices for an agent. If the actions and plans differ significantly from each other, it is probably justified to have two separate practices. The level of detail of the practices then mirrors the level of expertise of an agent in a certain area. The process of learning will influence the level of detail of the represented social practices: an agent may learn to differentiate between two contexts, and in the process split up an existing social practice in two new ones. With this in mind, the level of detail might even vary between areas for a single agent.

As this is mainly a problem that needs to be solved on a case-to-case basis during the design of a concrete implementation, I will not discuss it in the following sections. It should be kept in mind, though, that a form of social learning can be represented by the gradual refinement of social practices.

⁹It might be desirable to represent social practices in a hierarchy of increasing detail (see section 5). In this case the requirement of mutual exclusivity should be weakened: social practices that are each others direct ancestors or children may (must?) apply in the same context.

Perception	<i>Physical</i>	Resources	Actors	Places
⋮	<i>Social</i>	Uses	Roles	Interpretation
Recognition		Norms		
Application	<i>Actions</i>	Competences	Plan	
⋮		Activities	patterns	
Deliberation		Meaning		

Figure 2: Schematic representation of a model for social practices.

3.4 Modified model

Following the discussion of this sections, I would like to present a slightly modified version of the representational model by Dignum and Dignum [5]. This model is presented in fig. 2.

The concepts discussed in this section are placed at the right. The second column from the left serves to roughly divide the concepts into categories or conceptual layers. In the first column, some processes related to these layers are mentioned. These will be discussed in the next section.

The dashed lines between concepts indicate a close relationship. The physical concepts and their social counterparts share such a kinship, defined by a many-to-many relationship (as discussed above). The lines are dashed only when such a relationship is exclusive. Meaning, for example, has a close relationship to activities, but to plan patterns as well. Competences, in contrast, share a relationship exclusively with activities.

The layout of the *actions* category requires some additional explanation. Competences are closely linked to activities, hence their positioning in the same block. Plan patterns are closely related as well, as they provide structure and constraints for the organisation of actions. This relationship is not clearly shown in the model, but their positioning mirrors the closeness of the concepts. Meaning is related to both individual actions and plan patterns as a whole, and is thus a ‘mutual foundation’ in the model.

With this model in hand, I will set the first steps towards an architecture in the next section.

4 Applying concepts

In the previous section, I have discussed all concepts separately. In this one, I will play out the scenario discussed in section 2 for a police officer to gain an insight

in how these concepts fit together. In the process, additional requirements for an architecture built upon social practices will be discovered.

4.1 Working out the scenario

I will work out the scenario for a police officer. In the process, concepts that have not been precisely defined will surface. Those concepts will be in italics, and will be discussed in greater detail in the following sections.

Bob, a friendly policeman, is called upon by his chief to help maintain order during a small protest at a nearby car factory. When he arrives at the scene, he immediately *recognizes*₁ his colleagues and his commander. He *knows* that he should join them and he positions himself next to his colleagues. Bob *infers* that the people in front of them are the protesters, and he *recognizes*₂ that the situation is currently under control.

Suddenly, he *spots a change of the situation*, as a fight develops before him. He *decides* not to take action, yet. His commander then gives the order to break up the fight. Bob *obeys* the order, and ends the fight in *cooperation* with a few of his colleagues. Bob and his colleagues use as little aggression as possible, as they do not *want* to escalate the situation. Unfortunately, the situation turns sour, and soon Bob and his colleagues find themselves faced by an angry mob. They defend themselves, and they no longer hold back on aggressive actions. Several protesters are arrested, and eventually the crowd disperses. At this point, Bob and his colleagues regroup and eventually leave the location to transport the detained protesters to the police station.

4.2 Prerequisites

An agent needs certain components as a precondition for it to be able to use social practices. Nearly all concepts in italics in the scenario above have something to do with these prerequisites. They can roughly be classified in three broad categories: *perception*, *identity* and *deliberation*. I will briefly discuss each of them:

- *Perception* refers to the ability of the agent to perceive its environment and classify objects (and agents) in it in a certain way. This classification does not (yet) contain any social components: it should be based purely on the physical characteristics of objects in the environment. The classification can be as coarse or precise as is appropriate for the context, yet a more precise classification will probably lead to better (easier) matching with social practices.
- *Identity* concerns all the knowledge, beliefs, motives, goals, intentions, desires, etcetera (if any) other than the social practices. Identity may play a

role in decision making at all levels. For example: an agent might already *know* the role of a certain agent, aiding in recognition. Additionally, the *goals* an agent has, might influence plan selection within the practice and play a role further in deliberation after the social practice has been applied.

- *Deliberation* is the process of considering and selecting courses of action. The social practices do not replace this process, they should merely provide its starting point: a set of constraints on action sequences, from which a plan must be composed. The deliberation should be able to generate a plan, based on the input provided by the active social practice (and the identity of the agent).

4.3 Recognition

The process in between the input (perception) and the process where the output is used (deliberation), has yet to be defined. This is the place of social practices: they translate their input to sets of actions and/or constrains on action sequences.

Immediately following perception, *recognition* should take place. This term may be used for recognition on two different levels: recognition of objects, agents and locations (as having a certain role, use or interpretation), *recognition*₁, and recognition of context or situations, i.e. recognition of applicable social practices, *recognition*₂. Both types of recognition follow object recognition, which I have named *perception* in section 4.2.

Recognition₁ is unavoidably a part of (or precondition for) recognition₂. An agent first perceives the world around it, and forms a model of his environment. This is what I call *perception*: the forming of an internal representation of the environment on a basic level. This should be followed by recognition, i.e. the matching of available objects and agents to roles and uses (within a social practice). At this point, both recognition₁ and recognition₂ come into play. The matching of object to uses and agents to roles should trigger the best fitting context for the existing situation.

The translation is provided, as discussed in the previous section, by the physical layer. It should implement a translation both in terms of meaning and number.¹⁰

The problem of this translation has been considered before. In [10], a model is proposed based on dramaturgical theory. I will not go into the details here, but this approach is on many fronts similar to the one I have taken. The recognition process I discussed could be implemented in a similar way as is proposed in that paper: probabilities are assigned to role attributions (which in our model should be extended with use-attributions and interpretation-attributions for locations)

¹⁰By ‘meaning’ I refer to what agents *count as* an instance of a role. By ‘number’ I mean the (relative) quantity of agents and objects of certain roles or uses.

based on perception, historical attributions and expectations. Kochanowicz et al. treat scenes¹¹ separately: they are processed in parallel in a way similar to the processing of roles: based on expectations and perception a scene is recognized.

The last part is, at least in this case, undesirable. Scenes should not be considered separately. The model of social practices does not have a separate representation of situations or scenes. The social practice itself defines the context of actions, and can be seen as representing a situation. The issue of recognizing which social practices are applicable is the very issue that is under consideration here. In this context, there is no benefit in defining recognition of social practices in terms of probabilities based on perception: it merely restates the problem in different terms. The model of Kochanowicz et al. presumes recognition of high level concepts by the agent. I would like to avoid this where possible.

Instead, the recognition of a social practice should be broken down into recognition of agent roles, object uses and location interpretations. Previously, I mentioned that the physical layer should provide the link. This does not necessarily imply a one-to-one matching between perceived objects and possible uses. The idea of assigning probabilities, as is exemplified by the Kochanowicz's model, could fit well here. One might equip an agent with a database that assigns the probability of 0.8 for guns having 'enforcement equipment' as a use.¹² The idea of interpretations of objects, agents and locations reinforcing or inhibiting each other might be useful as well: the presence of reinforcement equipment increases the probability of the presence of police officers, for example.

A rather complex calculation is then required to compute the probability of role and use assignments. One should avoid making this process overly complex. If a complex calculation has to be made for every perceived object, also taking into account the probabilities of all other assignments, the time required for this grows exponentially with the number of perceived objects. This could be partially remedied by splitting the calculation in a two step process: in the first step, probabilities are calculated *an sich*, that is, for each individual object with no regard for correlations with other perceived objects. In the second step an update can be done using probabilities of simultaneous occurrence, using the computed values from the first step. This second step could be repeated multiple times, but I doubt that many repetitions will improve the recognition significantly. The costs will quickly outbalance the benefits.

A possibly suitable technique for the implementation of this 'dictionary' might be an artificial neural network. ANNs handle probabilities (weights) and interde-

¹¹Scenes in the sense that Kochanowicz et al. use the term, are holistic representations of situations.

¹²This value is rather different for different countries. In the Netherlands, for example, civilians are not usually permitted to carry a weapon, so a high probability is justified here. This would of course be different in the United States.

dependencies (links) fairly well. The use of ANNs might enable the merging of the perception and recognition processes. ANNs are already in use for (visual) object recognition.¹³ Of course, this depends entirely on the application. For relatively simple uses, ANNs are probably too heavy an instrument for this purpose.

Once the probabilities of assignment for individual objects have been established, the link with social practices should be made. Here, a choice presents itself. I see two possible ways of going about this process:

- Select the highest probability assignment for every object. Find the social practice that matches these assignments most, or
- Compute the resulting probability of every practice using the individual assignment probabilities, and select the practice with the highest probability.

Of these two options, the second seems the most intuitive. Yet the first is probably much more efficient. Ideally, a middle ground between these two extremes should be found. Again, a two step process might be used, where pre-selection is done based on the first approach – a rather coarse evaluation of the applicability of social practices. The second step would then evaluate the best fitting social practices (the ‘survivors’ from step one) in more detail.

The above constitutes a global idea of the direction an implementation of the perception–recognition process could take. It requires further elaboration and should be made more precise before it is suitable for implementation. This leaves room for further research. An schematic representation of the ideas provided in this section is presented in fig. 3.

An open question remains how often this process should be repeated, as it is rather (computationally) expensive and will not always provide new information. A partial answer to this might be the following: an agent will (re-)start this process when it:

- a) Runs out of actions, i.e. completed its (current) plan, or
- b) Encounters a change in the environment that is not foreseen in the current plan (e.g. a plan for *surveillance* of a demonstration might not provide actions to handle the situation when fights pop up, as this is the terrain of an entirely different practice), or
- c) Performs or perceives an action that (usually) triggers a new social practice.

This is only a global idea: further work on the exact functioning of this component is needed.

¹³For an overview of the applications of (deep) neural networks, including their use in object recognition, see [13].

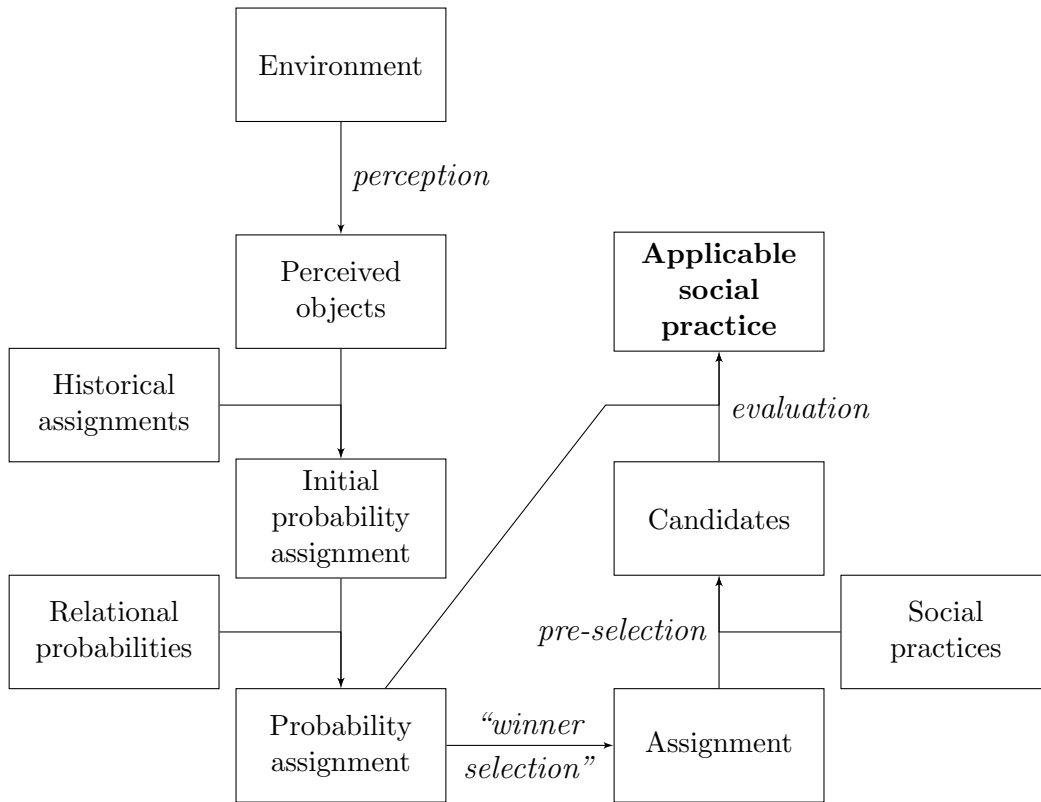


Figure 3: Schematic representation of the perception–recognition process.

4.4 Plan composition and action selection

I have sketched a way to “activate” a social practice. Once a certain social practices has been activated, all information contained within this practice becomes available. This should be supplied to the deliberation process in such a way that it may be used in this process. By providing constraints on actions through plan patterns, the basic social part of the decision process is contained within the social practice. It is left to the deliberation process to apply other considerations to the input provided by the social practice.

Once a social practice has been “activated”, the information an agent has about the world should be updated. In particular, I see three closely linked categories of information the agent should update:

- *Social model*: immediately following perception, an agent forms a representation of its environment (a *world model*). This models the physical aspects of the environment in a way that the agent understands. Similarly, an agent

should have a social model that represents social aspects of the environment. One could think of this as a *social landscape*, as described in [2]. I will not go into the details of the representation, but it is clear that the active social practice strongly influences this model. Relations among agents may change, and one might represent the current role assignment for agents in the environment within the social model, for example.

- *Identity* should contain all aspects that relate to the individual identity of the agent. Its beliefs, goals and general competences would belong here, for example. The current role of the agent itself within the active practice should be part of the identity as well. Although this may be redundant if it was represented already in the aforementioned social model, it belongs here conceptually.
- *History*. The agent should maintain some history of role assignment, object use, etcetera. Historical assignments may help in determining the correct assignments in the future (under the assumption that agents and objects keep the same role or use more often than not). This history can be as extensive or limited as desired: a more detailed history will lead to a more complex calculation of probabilities, and a larger administration, but might improve the accuracy of the selection process slightly.

As the definitive assignment of roles and uses becomes available only after a social practice has been selected, all this can only be updated after the selection process, but before (or possibly in parallel with) plan selection and execution.

With the social practice selected and all (social) information updated, we now arrive at the point where the deliberation process has to be provided with useful input. Several concepts play a role here: norms, competences, meaning, activities and plan patterns. That is, all the concepts linked with social practices, except the ones that were used in the selection of the social practice. Additionally, roles play a small part as well, in particular the role the agent itself has within the practice. I will discuss how all these concepts fit in the process.

Up till now I have assumed plan patterns are contained within the practice, but I have not discussed how an agent knows which plan pattern to select. Within many social practices, multiple plans will be available, as different roles will generally follow different plans. I would propose to have a plan available for each role in a practice, generally, no two roles will have the same plan pattern, but this is not impossible within this approach.

It is important to note that by role-specific plans, I mean the plans at the level of scene scripts. The general, high-level interaction structure that defines the relations between scenes is not role-specific.

This implies that once the agent knows its role, it knows which scene scripts apply. And that means that there can be only one plan pattern for each role. The question is if this is a realistic and viable solution. It eases the selection of a plan pattern, but it has a possible undesirable effect on social practices: once the plan pattern for one of the roles within a social practice changes slightly,¹⁴ a new social practice has to be formed. This might lead to a large number of practices.

I argue that this is both realistic and unproblematic. A social practice more or less mirrors a situation. If plan patterns differ, this is always due to a (slight) difference in the situation. It is then justified to create a separate practice: if an agent has *learned* to differentiate between two plan patterns/situations it has reached a level of expertise that warrants such further refinement. This is completely in line with the idea of the specificity of the social practices as mirroring the level of expertise, that I put forward earlier. Besides, how would an agent distinguish two plans if it fails to see the difference between the situations where they apply?

One could argue that this difference between plans is not necessarily of a physical nature. While this report has mainly been concerned with the physical aspects of the environment, it is not impossible to include emotional components (or other social elements) in the process, if the agent can recognize this. This might be handled within a social practice, for example, by assigning a mood to every agent and use this in plan selection. Social practices would then provide plans for every role-mood combination (probably not detailing a plan for every mood, but using *defaults* and more suited plans for particular emotions). As the social model already plays a role in the assignment process, such components could also enter the process there.

Lastly, one could say that it is unrealistic for an agent to only have one possible plan in a situation, however specific, as the actual action taken by an agent always depends on more than the (social/physical) situation in the environment. Although this is most certainly true, this goes beyond the scope of social practices. The actual decision making is not done within the social practice: this is left to the deliberation process. The social practice merely defines constraints on actions – or an action framework, if you will. It is up to the deliberation process to convert this to an actual action sequence, and take the agent’s identity into account in the process. A plan pattern may define two different possible courses of action: it only ensures that only one of those will be executed, and defines constraints on their applicability. In other words: choice is possible within plan patterns.

This leads the discussion towards the activities themselves: they are contained in the plan patterns, and are as such provided to the deliberation process. Ad-

¹⁴Change here is not meant as change by learning, but as the difference between two possible plan patterns.

ditionally, they carry some preconditions: competences. In my description of the agent identity above, I have placed the agent's competences in the identity. This would mean that the deliberation process would need to consider if the competences of the agent enable it to execute a particular activity. A point of critique here could be that not all competences are entirely individual, some come with a role (although these would be mostly competences in the sense that the agent is *allowed* to do something, arrest someone, for example). This is not a problem here, as a plan is already directly linked to a role. One can assume that a plan pattern will only be passed if the agent has (one of) the appropriate role(s), and thus has these competences. They do not need to be defined explicitly for each activity.

Norms require some attention as well: they are a difficult concept in this context. Norms are rules of expected behaviour. As far as this applies to the actions of the agent itself, norms do not add much to the mix: it is expected that the (expected) behaviour of the agent is captured in the plan patterns, which would render norms redundant in this case.

The function of norms as predicting the behaviour of others falls victim to the same argument, as plan patterns are available for all other roles, as well. Norms might serve as a short-cut though, as it is infeasible to run the deliberation process for all other roles (or agents, even).

But norms serve a different purpose. They define the boundaries of acceptable (expected) behaviour. Norms become important once they get violated. Plan patterns might define fairly strict usual courses of action, but norms define the broader limits of what is acceptable within the practice. If a plan is not followed, but no norm is violated, the agent cannot assume a context change (a change of social practice), but likely has to adapt its plan and adjust its expectations. If a norm *is* violated, the agent needs to re-evaluate which social practice applies.

Usually, some meaning is coupled with the violation of a norm. And thus, implicitly at least, with compliance as well. Norms are boundaries, and there are consequences attached to the crossing of those boundaries. Some issues coupled with the use of norms in multi-agent systems are summed up in [14].

Norms do not have an explicit place within the architecture I discussed. Further thought is required on this subject.

The last concept to be considered is meaning. Meaning should be represented as explicit changes to the social model that mirror the social effects of an action. Earlier I said that meaning is attached to both activities and plan patterns. It would make sense to attach meaning (explicitly) to activities and sequences of activities within a plan pattern. When an activity is then actually selected for execution, or a sequence of activities has been completed, the change to the social model is triggered. Providing meaning with the activities also leaves room for the deliberation process to reason about this, although this is not a requirement

for the deliberation process. One might expect social agents to have social goals, besides the traditional goals (in their identity). When faced with a choice within a plan pattern, an agent might reason about which activity (sequence) brings it the closest to fulfilment of that goal.

In conclusion: the social practice severely restricts the options for the deliberation process to choose from, which is exactly what we want. This comes at a cost, a part of which is that the deliberation process has to meet certain requirements, namely:

- It should be able to handle plan patterns, and generate plans from it. How it does this exactly is left, mostly, for the implementer to decide.
- It should check the competences of the agent before selecting an action.
- It might reason about meaning and social goals.

Next I will try to integrate the considerations from this section in a basic architecture for agents using social practices as the basis for their deliberation.

5 Architecture

In the previous sections, I have extensively discussed the concepts and processes involved in an implementation of social practices. The first steps to a more concrete architecture have already been made in the previous section. In this section, I will integrate the work from the previous sections and present a (partial) architecture for the implementation of social practices.

5.1 Abstract architecture

In the previous sections, I have shown that social practices can reduce the number of applicable actions immensely. If an efficient implementation for social practices within an agent architecture can be provided, this can potentially speed up the deliberation process significantly. When the context has triggered a social practice, and this leads to the availability of only one possible action to the agent, the remainder of the deliberation can be skipped altogether. When multiple actions remain, the social practice will provide aid in limiting the options that are to be considered.

An abstract agent architecture that makes use of this principle was presented in [6] (fig. 4). It places the social practices at the basis. The context triggers one or more social practices, and one is selected taking into account the agent's identity. This social practice then provides all the tools for the deliberation process to generate plans, again, taking into account the identity of the agent. Once a

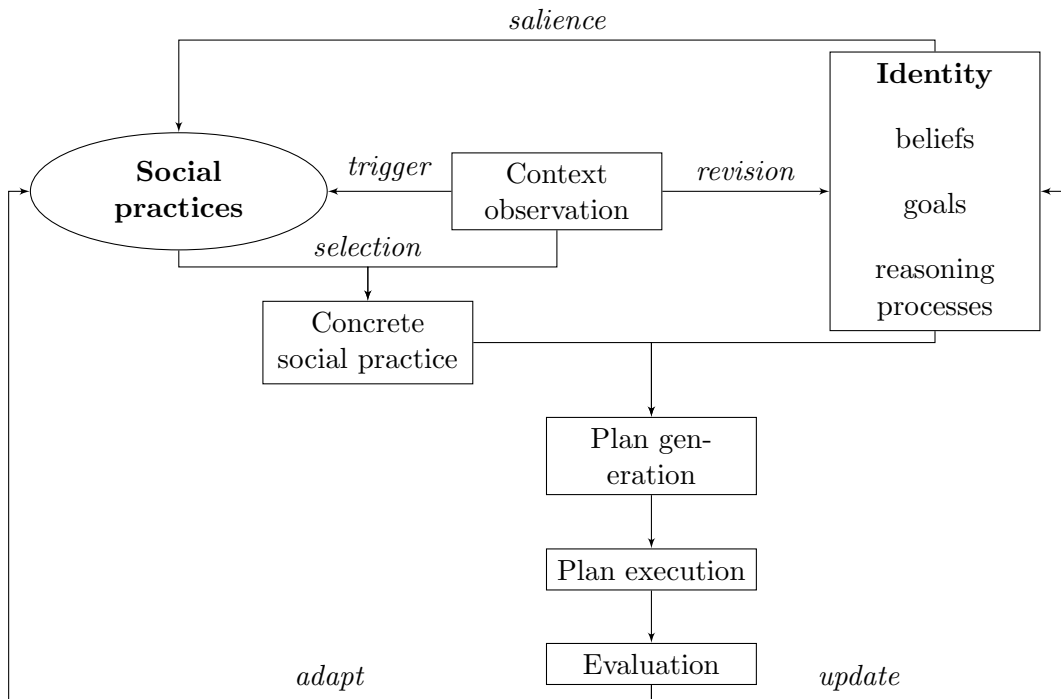


Figure 4: Abstract architecture proposed in [6], figure from [5].

plan has been selected and executed, learning may take place while evaluating the results. This may lead to the adjustment of both the social practices and the identity of the agent.

Many of the processes in fig. 4 share a similarity with processes that have been discussed in previous sections. The processes represented by the arrows labeled *Trigger* and *selection* have been worked out in (relative) detail, together with the output from the social practice towards the deliberation process. The *revision* and *update* of the identity has been touched upon, as has the *salience* that the identity would provide for social practices. The *adaptation* of social practices has been hardly discussed in this report, although I mentioned it on a few occasions when referring to learning.

I have exploited these similarities, and have tried to come up with a merge of this model and the bits and parts I assembled in this paper. The result is shown in fig. 5.

The architecture needs some additional explanation. The main flow from perception to execution (input-to-output) is shown by the solid lines. Side processes and side effects have been drawn with dotted lines. I do not claim this architecture to be complete, but it represents the ideas that have been discussed in these reports. It cannot be seen separate from the text though, as the descriptions in

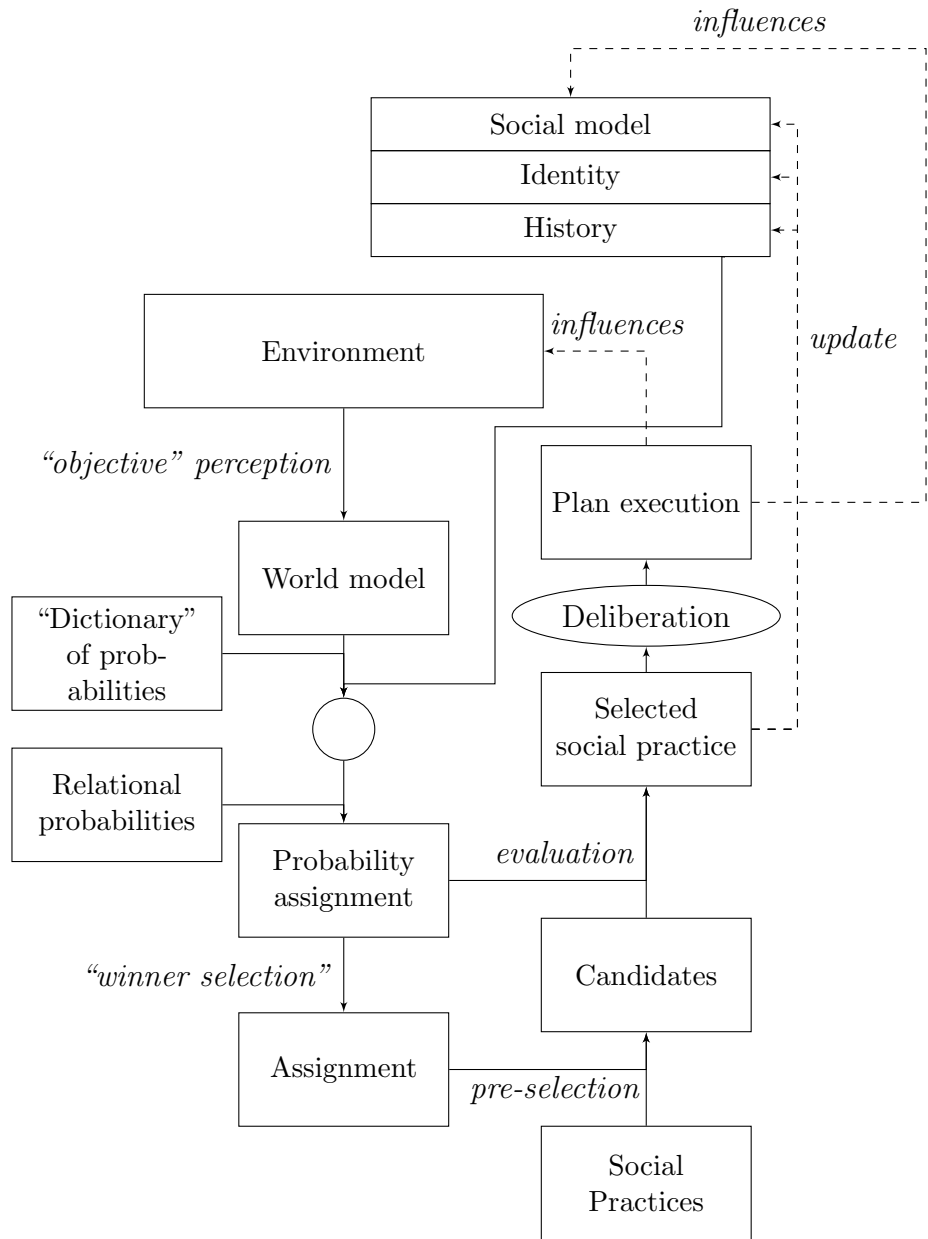


Figure 5: Base architecture for an implementation of social practices.

the previous sections provide details that are vital to the understanding of this figure.

5.2 Loose ends

With regard to the architecture in fig. 5, some remarks have yet to be made that were not covered in the previous sections. I will highlight them here.

I have not gone into the implementation of the several different storages: not of the probability “dictionaries”, not of the social model, identity and history and not of the social practices itself. I will only say something of the last one. The exact implementation, beyond any constraints that have been exemplified in previous sections, is outside the scope of this report.

The structuring of the database of social practices is a different matter. The selection of social practices takes place on the level of actors and resources, i.e. on the physical level. Any structure that is used for the storage of social practices should be indexed on those concepts, as it should be easily searched during the selection process.

Preferably a hierarchy of practices should be constructed, with more abstract practices at the top, and more concrete practices at the bottom. An agent can then go as deep in the structure as it has learned to differentiate between situations. This would eliminate some redundancy, as it is not required to test requirements multiple times: if the parent satisfied them, the child should as well. The exact way that this structure should be built up is not yet clear, and this requires some further consideration. An example could be taken from [1], although rigorously applying the solutions provided there for solving overlap and sub-contexts might lead to unnecessarily complex inheritance structures. They should only be applied when sufficient similarity exists. Additional measures for reducing redundancy may be taken in the way the elements contained in social practices are stored.

The description of roles, actions, etc. should all be provided separately from the social practice. The social practice should as much as possible provide links to the relevant representations, unless it concerns content that is unique to the social practice, in particular the physical (translation) level. This minimizes redundancy and therefore increases efficiency.

Lastly, I will note that I have described a stimulus-driven approach in this report. In the proposed architecture, I did not take goal-driven selection of social practices into account. I will briefly discuss this here.

An agent is not always passively looking for the right social practice to apply. Sometimes, an agent might seek to initiate a certain practice, or expect a social practice to be active. The first case is not really different from the second. In this case, the agent will first undertake some action(s) to initiate the desired practice (possibly to fulfil one of its goals). Once this has been done, it expects that practice

to be active.

An agent might expect a social practice to be active when it or other agents performed certain actions that would in the current context normally lead to the activation of that social practice. Another possibility might be that the agent waits for a certain situation to occur, as someone at a bus stop waits for the bus to arrive.

Take for example Bob, the policeman from section 4.1. When he arrives at the scene, he probably expects the *small-scale-protest* social practice to apply before he has observed the environment, as he has been sent there by his chief to maintain order at a small scale protest. He does not need the environment to tell him what social practice applies, but to confirm that the *small-scale-protest* social practice applies.

In general, it means that the agent is actively trying to apply a particular social practice. The agent knows what it is looking for and because of this foreknowledge there is no need to start searching bottom-up with the elements in the environment. Instead, the agent can take a short-cut. It knows what practice it wants to apply, it only needs to know if it can.

The practice provide the roles and uses that are expected. The processes at the bottom of fig. 5 can be skipped altogether. If the “*dictionary*” of probabilities is implemented as a two-way dictionary, it is imaginable that this is used to actively match elements in the environment with roles and uses. The probability assigned to the match can in this case serve as a measure for the certainty with which the match was made. If all elements of the social practice are matched with sufficient certainty, the social practice is applicable.

This short track to social practice recognition is currently not integrated in the architecture. I leave it to future work to work out the details and integrate this track in the process.

6 Conclusions

In this report I have discussed the theory of social practices and its application in agent technology. Taking the representational model of Dignum and Dignum [5] as a starting point, I extensively discussed the concepts contained in their model, and proposed a model of my own based on this discussion.

I have extracted two important processes in the application of social practices: the activation of the social practice (perception – recognition) and its use (application – deliberation). Based on a discussion of these processes and the abstract architecture from [6], I have come up with a proposal for a base architecture (fig. 5).

I have shown that it is, indeed, possible to provide a precise description of

the concepts related to social practices: I have been able to transform the vague sociological concepts into concrete notions that are ready to be tested in an actual implementation. In the process of making these concepts concrete, I have provided an overview of the issues related to implementing social practices for agent systems. Despite the elaborate discussion, the architecture I described is not yet complete, nor perfect. It may be further extended and elaborated, but I have taken a large step towards an actual implementation.

Such an implementation will enable agents to exhibit increasingly social behaviour and reason about social consequences. By using social practices as the basis of the deliberation process, the reasoning of the agents is inherently social. Artificial agents will be better able to interact with and understand human agents, and vice versa. They will be better able to determine appropriate behaviour in certain contexts, comply with “unwritten rules” and so forth.

Before this is realized, additional work on the topic is required. This report did not cover all aspects of an implementation. I will list some suggestions for building upon this report below:

- The learning/adaptation process for social practices has to be worked out.
- The abstract process descriptions given in section 4 should be elaborated, evaluated and preferably implemented and tested.
- The edge-components need to be worked out: the social model, identity and history.
- The short track of social practice recognition should be integrated in the architecture.
- Stricter constraints on the rest of the agent architecture (world model, deliberation, etc.) should be defined, and possible implementations need to be considered.
- Further work is required on the place of norms in this architecture.
- Generally, implementations of agents using social practices need to be developed and tested.

Although the list above poses a lot of challenges before the truly social artificial agent is a reality, the first steps have been taken towards an implementation of social practices for agent systems. I hope to have aroused thought about the concepts of social practices in my detailed discussion of them, leading to the representational model of section 4. Also, I hope that my considerations towards an architecture for the social agent will set the development of such an agent in motion. That we may interact with our artificial friends in an intuitive way in the near future!

References

- [1] H. M. Aldewereld, S. Alvarez-Napagao, F. P. M. Dignum, and J. Vázquez-Salceda. Engineering Social Reality with Inheritance Relations. In H. M. Aldewereld, M. V. Dignum, and G. Picard, editors, *Proceedings of the 10th International Workshop “Engineering Societies in the Agents’ World” (ESAW 2009)*, volume LNAI 5881, pages 116–131. Springer Berlin Heidelberg, 2009.
- [2] F. P. M. Dignum, R. Prada, and G. J. Hofstede. From Autistic to Social Agents. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Paris, France, May 5-9, 2014*, pages 1161–1164. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [3] M. V. Dignum. *A Model for Organizational Interaction: Based on Agents, Founded in Logic*. PhD dissertation, Utrecht University, 2004.
- [4] M. V. Dignum and H. M. Aldewereld. OperettA: Organization-Oriented Development Environment. In M. Dastani, A. El Fallah-Seghrouchni, J. F. Hübner, and J. a. Leite, editors, *Languages, Methodologies, and Development Tools for Multi-Agent Systems*, pages 1–18. Springer Berlin Heidelberg, 2011.
- [5] M. V. Dignum and F. P. M. Dignum. *Contextualized Planning Using Social Practices*. Unpublished manuscript, 2014.
- [6] M. V. Dignum, C. M. Jonker, R. Prada, and F. P. M. Dignum. Situational Deliberation Getting to Social Intelligence. In *Computational Social Science and Social Computer Science: Two Sides of the Same Coin, Guildford, England, 23-24 June, 2014*, 2014.
- [7] S. Hofmann. *Dynamik Sozialer Praktiken: Simulation Gemeinsamer Unternehmungen von Frauengruppen*. Springer-Verlag, 2009.
- [8] G. Holtz. An Agent-Based Model of Social Practices. *ESSA 2012 Proceedings*, pages 1–10, 2012.
- [9] G. Holtz. Generating Social Practices. *Journal of Artificial Societies and Social Simulation*, 17(1):17, 2014.
- [10] J. Kochanowicz, A.-H. Tan, and D. Thalmann. Dramaturgical and dissonance theories in explicit social context modeling for complex agents. *Computer Animation and Virtual Worlds*, 26:247–257, 2015.

- [11] J.-J. C. Meyer and R. J. Wieringa. *Deontic Logic in Computer Science: Normative System Specification*. John Wiley and Sons, Inc., New York, 1994.
- [12] A. Reckwitz. Toward a Theory of Social Practices: A Development in Culturalist Theorizing. *European Journal of Social Theory*, 5(2):243–263, 2002.
- [13] J. Schmidhuber. Deep Learning in Neural Networks: An Overview. *Neural Networks*, 61:85–117, 2015.
- [14] L. C. B. C. Vanhée, H. M. Aldewereld, and F. P. M. Dignum. Implementing norms? In J. F. Hübner, J. M. Petit, and E. Suzuki, editors, *Proceedings of the 2011 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology - Workshops, WI-IAT 2011, Campus Scientifique de la Doua, Lyon, France, August 22-27, 2011*, volume 3, pages 13–16. IEEE Computer Society, 2011.