



Modelling the Social Practices of an Emergency Room to Ensure Staff and Patient Wellbeing

Mercuur, Rijk; Larsen, John Bruntse; Dignum, Virginia

Published in:
BNAIC 2018 Preproceedings

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Mercuur, R., Larsen, J. B., & Dignum, V. (2018). Modelling the Social Practices of an Emergency Room to Ensure Staff and Patient Wellbeing. In *BNAIC 2018 Preproceedings* (pp. 133-147)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Modelling the Social Practices of an Emergency Room to Ensure Staff and Patient Wellbeing

Rijk Mercurur¹[0000-0002-3956-2284], John Bruntse Larsen²[0000-0003-2500-9944],
and Virginia Dignum¹[0000-0001-7409-5813]

¹ Delft University of Technology, Postbus 5, 2600 AA Delft, The Netherlands
{R.A.Mercurur,M.V.Dignum}@tudelft.nl

² DTU Compute, AlgoLoG, Technical University of Denmark, 2800 Kongens Lyngby,
Denmark jobla@dtu.dk

Abstract. Understanding the impact of activities is important for emergency rooms (ER) to ensure patient wellbeing and staff satisfaction. An ER is a complex social multi-agent system where staff members should understand the needs of patients, what their colleagues expect of them and how the treatment usually goes about. Decision support tools can contribute to this understanding as they can better manage complex systems and give insight into possible problems using formal methods. Social practices aim to capture this social dimension by focussing on the shared routines in a system, such as diagnosing or treating the patient. This paper uses the Web Ontology Language (OWL) to formalize social practices and then applies it to the ER domain. This results in an ontology that can be used as a basis for decision support tools based on formal reasoning, which we demonstrate by verifying a number of properties for our use case. These results also serve as an example for formalizing the social dimension of multi-agent systems in other domains.

Keywords: Architectures for social reasoning · Ontologies for agents · Cognitive models · Agent-based analysis of human interactions

1 Introduction

A better understanding of the impact of activities in an emergency room (ER) on patients and staff would improve the wellbeing for both of them [12]. A decision support tool that has knowledge about ongoing activities could assist ER management in directing interventions by identifying what activities take place, what causes them to take place and who are typically involved in the activities. An ER is a complex social multi-agent system, where staff members should understand the needs of patients, what their colleagues expect of them and how the treatment normally goes about [20]. Focussing on this social dimension in these decision support tools could increase their realism [13] and therefore their potential in giving helpful insights into the domain.

So far ER models have mainly focussed on the patient flow, but not the interaction of the staff. An ER is modelled from a control flow perspective based on

explicit regulations and clinical guidelines [16]. This is useful to identify possible process bottlenecks, but offers little insight into their social causes. A possible cause for the lack of sociality in decision support tools for ER might be that there is no formal model that expresses this social dimension.

To use a formal reasoner to identify social bottlenecks three steps are involved (1) fairly representing the social dimension of an ER in a model (2) formalizing the model and using a formal reasoner to check if the model satisfies certain properties (3) translating these results back to interventions in the real world ER. This paper focusses on the second step by defining a precise, unambiguous and consistent semantics to support decision support tools. In particular, we aim to give a basis for formal reasoners that can infer helpful new social knowledge about an ER. Building upon [15, 19, 7], we use social practices to capture the social dimension of a system. The social practices in a system are the routinised actions that are (to some extent) similar for all agents such as, diagnosing or treating the patient. We use the Web Ontology Language (OWL) [21, 4, 10, 9], based on Description Logic, and the Protégé tool to formalize a social practice meta-model and apply it to the ER domain. This results in a formal ontology that includes a social practice meta-model and a domain specific ER model that can be used for decision support tools. We demonstrate, with data based on a visit to Herlev Hospital in Denmark, that we can express and verify whether a number of social properties that one could desire of an ER holds true in our model. This serves as a proof of concept to show that our formalization can be used to infer helpful new knowledge about a system. These results serve not only as a first step to better decision support tools for ER, but also serves as an example for formalizing the social dimension of other multi-agent systems.

Section 2 presents some background knowledge on social practices. Section 3 introduces the social dimension of the ER domain, its empirical grounding, and lists a number of exemplary social properties one could desire of an ER. Section 4 formalizes the social practice meta-model and our ER use case. Section 5 formalizes the earlier stated social properties and uses a formal reasoner to verify them (i.e., infer new knowledge about the system.) We end the paper by showing how our formalization is unique in this ability to capture social properties by comparing it to other agent meta-models.

2 Background

The concept of social practices stems from sociology, and aims to depict people's 'doings and sayings' [18, p. 86], such as dining, commuting and greeting. [19] recently revived the concept for its ability to highlight that our actions can be captured in routines that are similar for many people. For example, doctors and nurses follow similar routines in treating a patient. [7] proposed to use social practices to capture the social dimension of agent decision-making. They connected the concept to more standard agent concepts, such as actions, plans and norms. This section introduces a meta-model for a social practice agent that is expressed in the Unified Modelling Language (UML). The remainder of this

paper focuses on the formalization of this UML-diagram in OWL. This formalization provides an unambiguous basis for decision support tools in social systems such as the ER. Future work will focus on how the UML-diagram is grounded in the social practice literature and the justification of the model choices.

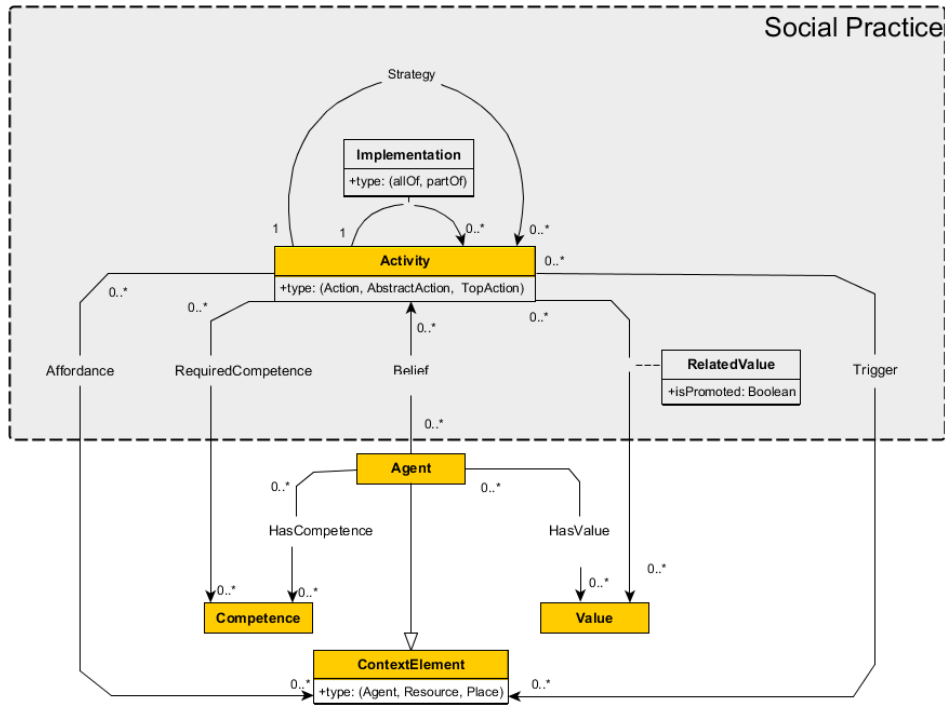


Fig. 1. The social practice meta-model captured in the Unified Modelling Language, including classes (yellow boxes), associations (lines), association classes (transparent boxes), navigability (arrow-ends) and multiplicity (numbers).

Figure 1 shows the social practice meta-model in a UML-diagram [17]. The main classes for a social practice model are activities (e.g., diagnosing, assign team), agents (e.g., nurses, doctors), competences (e.g., perform triage, managing team), context elements (e.g., phone, bed, triage room, other nurses) and values. Values here refer to human values, that is, ‘what one finds important in life’ [14], such as health or education. The social practice is an interconnection of (1) activities and (2) related associations as depicted by the grey box in Figure 1. For example, the social practice of acute treatment consists of several activities, such as assigning teams or performing triage. Figure 2 shows all the activities that the social practice of acute treatment comprises. Section 3 will further explain what these mean, for now it is important to understand the structure of this

activity tree. Shallow nodes are more abstract activities (e.g., assign to team) and deeper nodes are more concrete ones (e.g., inform patient). The type of an activity (respectively, **AbstractAction** and **Action**) is captured in the **type** attribute in the UML. The social practice connects these different activities with the **Implementation** association. If activity *A* implements activity *B* this means that *A* is a way of or a part of doing *B*.

Implementation is the first of several activity-associations specified in Table 1. Most associations are fairly self-explanatory, however the **Trigger** and **Strategy** association are a bit more complex. Following [23], triggers are the basis for habitual behaviour. If an agent is near a context element that has a trigger association with an activity, then it will do that activity automatically (without for example considering its values). Following [5], strategies are the basis for expectation about the actions of others. If an agent believes that (the personal part of) activity *A* is a strategy for activity *B*, then it believes that other agents usually implement activity *B* by doing activity *A*. Fundamental to social practices are that many of these activity associations are similar (or: shared) for most agents. Section 4.2 will explain how we capture this similarity in the model.

Table 1. The associations attached to an activity and their specification.

Association	Specification
Implementation	which activities are a way of or a part of doing the activity
Affordance	which context elements are needed to do the activity
RequiredCompetence	which competences are needed to do the activity
Belief	which activities an agent has beliefs about
RelatedValue	which values are promoted or demoted by the activity
Trigger	which context elements habitually start the activity
Strategy	which activities usually implement the activity

The agent furthermore has two associations, which plays a role in choosing the activities it will do: **HasCompetence** and **HasValue**. The **HasCompetence** association links possible skills to the agent who masters those. The **HasValue** association captures if an agent finds that value important.

3 ER Use Case

3.1 ER Description

An ER acts as the entry point for most hospitals and a wide variety of patients arrive there on short notice. The purpose of an ER is to provide immediate treatment for the patient and identify the further course of action. The general process of patient treatment can be seen as consisting of the following phases:

1. Contact ER: The secretary registers the patient who contacts the ER department.

2. Triage: A nurse performs triage on the patient in order to identify how urgent or life critical the patient is.
3. Tests: Doctors and nurses perform tests on the patients.
4. Diagnosis: Doctors give a (partial) diagnosis and perform treatment.
5. Plan: Doctors, together with the patient, establish a plan for further course of action.

Staff members thus comprise doctors, nurses, but also persons who are in charge of administrative tasks and of being in contact with patients. In some cases, an ER also serves as a learning facility for healthcare staff in training. It is common to have trainees participate in the treatment to get work experience as part of their education. Some trainees may have enough experience to carry out tasks by themselves, while others must be accompanied by more experienced staff.

The staff thus needs to cooperate and coordinate their actions, while being flexible enough to handle a wide range of patients. This means that staff members should understand the needs of other staff members and have an idea of what guides their actions. A staff member should be aware of the culture of the organization by understanding what is important and how things normally go about. He or she needs to know what is expected of him and what not. All of this facilitates the teamwork that is needed for a properly running an ER. As should be evident from the above description, social interaction is an important part of an ER, which makes it an interesting domain to apply social practices.

3.2 Empirical Grounding

The ER description is partially based on observations from a half-day tour at the ER department at Herlev Hospital in Denmark. The tour was led by head nurses from the department and consisted of a visit to the main reception, the trauma reception, and the areas of the three specialist teams of the ER department. During the tour the head nurses explained the rooms in the department, the organization, general work procedures, and how a typical work day went like. After the tour we observed some of the staff members in the specialist teams for a few hours, asking a few questions whenever possible. During the visit several staff members relayed some personal views on their work, including what they considered important in doing it. Note that to our knowledge there are no empirical studies of the social practices in an ER (e.g., [3, 1, 22] only provide high-level analytical statements), therefore our work is based on observations from this half-day tour. To go beyond the proof of concept presented in this paper a more extensive empirical study is needed.

Summary of relevant observations The general attitude in the ER department is geared towards being flexible and accommodating the needs of the patients. The established patient treatment procedures and task distributions are suitable for the treatment of most patients. However, the staff is open towards making adjustments if they believe the adjustments can help. The following list comprises some examples of social intelligence we observed:

- Sending samples to analysis is supposed to be done by nurses, but because many patients required attention, the head nurse helped with sending samples to analysis to give the nurses more time to attending patients.
- The secretary is supposed to receive patient transports but the head nurse helped out with this as well sometimes, leaving a note that the secretary would later then register in the system.
- In the middle of the staff room there was a box with current tasks. Nurse students carried out some of the tasks under supervision of nurses who assigned the tasks to the students. When the number of tasks had grown large, the head nurse called one of the other head nurses to ask if they could help with some of the tasks.
- The secretary assisted with attending patients to the extent possible.

3.3 Desirable Social Properties

To evaluate if we can indeed express and reason about the social dimension of the ER, we state a number of concrete social properties one could desire, which reflects some of the observations we made during the tour:

1. The staff understand the needs of the patients.
2. A head nurse can cover some of the necessary tasks of the secretary.
3. The staff can help each other out, because they know the equipment the others need.
4. Nurses should follow directions from the head nurse.
5. The length of stay of patients is not longer than needed.

To give an idea of the scope to which our proof-of-concept can be generalized we aimed at a diverse set that relates to both patient and staff, both ambiguous and unambiguous and general and domain-specific statements. To make claims about how well our model could capture any social property lies outside the scope of this paper. Section 5 shows that we can express the first three complex social statements with our formalization and verify their truth using a formal reasoner.

3.4 Social Building Blocks

The core of the ER model will be the activity tree that consists of instantiations of the `Activity` class and the `Implementation` relation between them (see Figure 2). The activity tree roughly shows the phases of acute treatment represented as social practice activities. The root of the tree represents the top action, the leafs represent the actions and the nodes in between represent abstract actions. It breaks the Acute Treatment activity down into more concrete activities to match the phases of patient treatment in ER. Note that the activity tree includes administrative activities such as assigning a patient to a specialized team within ER. We have included administrative activities to show how social practices can provide insight into these activities as well as activities that involve the patient.

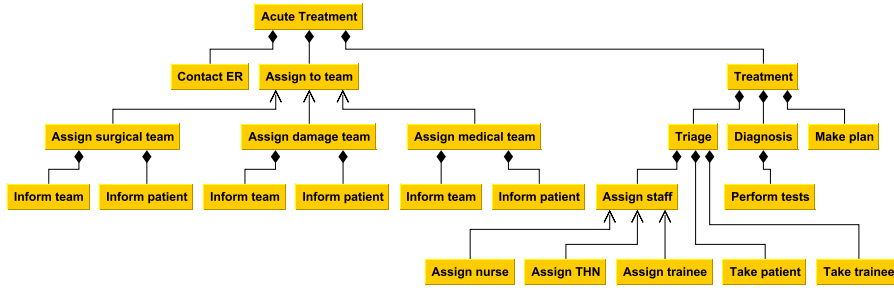


Fig. 2. The activity tree that (roughly) represents the different phases of ER patient treatment. The arrows with a diamond end represent partOf-implementations and the arrows with an arrow end represent allOf-implementations.

The ontology is based on this activity tree, but also provides more details of social practices in line with the UML of Figure 1. Table 2 provides an overview of agents, values and competences in the ER case. We only consider a small number of values and competences from the use case in order to make it clear how they are represented and reasoned with.

Table 2. Overview of agents, the values they adhere to and their competences.

Agent	Patient	Secretary	Head nurse	Nurse	Exp. Trainee	New Trainee	Doctor
Values	Health	Prof.	Prof., Educ.	Prof.	Educ.	Educ.	Prof.
Competences	Feedback	IT	IT, team	Triage, social	Triage, social	Social	Medical

4 Formalization

4.1 Meta Model

This subsection explains how the social practice meta-model described in Section 2 is formally captured in the Protégé tool. This comprises (1) translating UML concepts into equivalent assertions in the OWL syntax, (2) solving ambiguity and (3) making additional assertions that could not be captured in the UML syntax. The full formalization expressed in OWL can be found on GitHub.³

Class Hierarchy The UML classes can be translated one to one to OWL classes. However, we also need to specify the relation between classes. We can

³ The full formalization can be found on <https://github.com/PCSan/SOPRA>.

specify that an individual can not be a member of two different classes with the disjoint class axiom. Note that the `Agent` class and the `ContextElement` class are not disjoint. The UML shows this with the generalization arrow that goes from `Agent` to `ContextElement`. Following [2] such an association can be captured in OWL with the subclass axiom. In addition, we choose to capture the different activity types (i.e., `Action`, `AbstractAction` and `TopAction`) and different context elements (i.e., `Agent`, `Resource`, `Place`) as separate classes instead of attributes. This choice fits well with the notion of classes and avoids introducing an auxiliary typing scheme just for activities. This thus introduces three new disjoint classes (`Action`, `AbstractAction` and `TopAction`), which are all subclasses of the `Activity` class. This class correspondence can be captured in the ‘disjoint union of’ axiom.

Class Associations Each association is translated to an object property in OWL (except for the earlier discussed ‘generalization’). Object properties are defined by which two classes they connect (called respectively domain and range). We want to highlight a few interesting modelling choices. First, given the open world assumption of OWL⁴, we do not need to explicitly assert statements in OWL about the possibility of any number of associations (like is done in UML with the ‘0...*’ symbol). However, the ‘1’ multiplicity, that restricts the amount of classes does need to be asserted. We simply capture this in an object property cardinality restriction. These restrictions are not expressed as isolated statements, but instead as a restriction on a class. For example, we restrict the `SharedPart` class by saying it is a subclass of all entities that implement exactly one activity. Second, some of the associations in the UML have attributes. Given that OWL supports the notion of a subproperty, we capture these attributes in the ‘disjoint union of’ axiom (analogue to the `type` attribute of the `Activity` class).

Complex Rules Using SWRL⁵ we can assert a number of propositions that correspond to the semantics of social practices, but could not be expressed in UML. Not only are these SWRL rules crucial to making inferences about social properties, they also depict new insights we gained due to formalizing the UML in the OWL language. First, it allows us to assert several propositions regarding the relation of `Action`, `AbstractAction` and `TopAction`. Top actions are defined as activities that implement nothing, whereas actions are defined as activities that are implemented by nothing. Lastly, abstract actions can be found in the middle of an activity tree, implementing some other activities and being implemented by some other activities. Second, we can sometimes infer that if an object property relates to one activity (e.g., `affordance`, `relatedValue`, `requiredCompetence`), then it should also relate to another activity. For example, if `TriageExperience` is a `requiredCompetence` for `Triage`, then it should

⁴ The OWL reasoner assumes that as long as we do not explicitly state something *not* to be the case, it is possible.

⁵ SWRL is a language for extending an OWL ontology with Horn-like rules [10].

clearly also be a `requiredCompetence` for the `AssignStaff` activity, as this is a part of `Triage`. In other words, given that `AssignStaff` implements `Triage` it inherits some of its object properties. Not only the competences of a parent activity are inherited, but also its related values and affordances. For example, if the context element `ER-room` affords `AcuteTreatment` it also affords `Triage`. Third, as mentioned a strategy captures which activities usually implement an activity. To be more exact, a strategy connects the activity *A* and activity *B*, if an agent believes that activity *A* is usually done as an implementation of *B*. Thus for a strategy to exist between *A* and *B*, *A* at least has to be an implementation of *B*.

4.2 Model

Having a formal meta-model for social practices, we use it to formalize a model for the ER use case. The model asserts facts that are roughly based on our half-day tour at Herlev Hospital in Denmark. In a later section, we use formal reasoning to verify properties that follow from these assumptions. Note that to go beyond the proof of concept presented in this paper a more extensive empirical study is needed.

Shared and Personal Activities Social practices capture the social world in terms of shared activity associations. This means agents can have the same beliefs about activities as other agents. In our model, for example, all agents believe that the activity `Triage` requires the `Triage` competence. However, there are also personal views on activities. For example, a patient might believe `Triage` only relates to the value of `Health`, while a trainee thinks it also relates to the value of `Education`. This is reflected in the model by having multiple instantiations of an activity: some shared and some personal. For example, there is one instantiation called `SharedTriage`, which all agents believe. Shared associations such as that the `Triage` activity requires the `Triage` competence are associated to this instantiation. There is also an instantiation called `TraineeTriage`, which is only believed by trainees. The personal view that the value of `Education` is associated with `Triage` is attached to this instantiation. The fact that agents can assume some beliefs are similar for other agents is what makes it that we can infer *social* knowledge (as we will show in Section 5) and is what makes our ontology a unique tool (as we will discuss in Section 6).

Activity Tree The parts of the meta-model that constitute the activity tree are the classes `Activity`, `TopAction`, `AbstractAction`, `Action`, and `SharedPart`, and the object properties `implementationPartOf` and `implementationAllOf`. We translate each box in the activity tree into the corresponding activity so that the root is an individual of `TopAction`, the leaves are individuals of `Action`, and all other boxes are individuals of `AbstractAction`. Initially we create one shared activity for each box, and later we create personal activities for the agents.

Agents and Strategies The agents represent persons in an ER such as a nurse or a patient. For each agent we define a named individual of the `Agent` class. We assert that these agents know about an activity using the `belief` object property. In addition, we use the `strategy` object property to assert which activities agents believe as common ways of doing things.

Competences, Values and Context Elements For each activity, we assert agent competences that are required for that activity, and which agents own those competences. For example, we assert that it is commonly believed that triage requires triaging competences by asserting `Triage` as a required competence of `SharedTriage`. We then assert that a nurse has the triage competence by using the `HasCompetence` object property. We also assert values that the agents adhere to and associate with the different activities. For example, we assert that it is commonly believed that triage promotes health, by asserting `Health` as a promoted value of `Triage`. Finally, we assert context elements that affords activities taking place, and which can trigger habits of some of the agents. For example, we assert that the ER facility affords contacting ER by asserting that `ER` is an affordance of `SharedContact_ER` and that taking contact in the ER is a habit of the patient by asserting it a trigger of `PatientContact_ER`.

5 Formalization & Evaluation of Properties

This section first formalizes the properties introduced in Section 3.3. The formalization follows three steps (1) expressing the statement in the more precise social practice terminology, (2) rephrasing it into a query and (3) rewriting it in the formal SPARQL syntax. We need to rephrase the properties into queries as Protégé can not evaluate the truth of a proposition, but rather gives back the lists of individuals that satisfy the query. In addition, this sections aims to verify the truth of the propositions, by evaluating what the query returns. We aim to demonstrate that our formalization allows us to reason about knowledge we did not assert. We leave claims about actual specific ER departments for future work.

Property 1 We make the property more precise by specifying what it means to ‘need’ and what it means to ‘understand’. We specify the ‘needs’ of a patient as those context elements and competences that are required for (or afford) those activities the patient finds important. Important activities are those that promote values the patient adheres to. We translate ‘understanding’ something to having beliefs about something. In this case, the staff thus needs to have beliefs about certain competences and context elements that are important to the patient. Note that if in our formalism an agent have beliefs about an activity, this implies it believes the related competences and affordances. We can thus specify this property as ‘All the activities that a patient believes promote values that are adhered to by the (same) patient, the staff have beliefs about’. We

can query which agents have beliefs about the values that are important to the patient and check to what extent the staff members are part of this list.

```
SELECT ?agent ?value ?activity
WHERE {
    ?agent a:belief ?activity.
    a:Patient a:HasValue ?value.
    ?activity a:promotedValue ?value}
```

The formal reasoner returns a list of important activities and agents that have beliefs about these activities. In our proof of concept, we find that for every activity important to the patient, there is at least one staff member that has beliefs about this activity. However, not *all* the staff members have beliefs about the needs of the patient. For example, contacting the ER department is important for the patient, because it promotes the patient's health. One of the needs of the patient is thus the context element 'phone' that affords contacting the ER department. However, the formal reasoner shows that in our use case only the overall head nurse is aware of this.

Property 2 We make the property more precise by specifying what 'can cover', 'necessary tasks' and 'tasks of the secretary' mean. We specify the 'tasks of the secretary' as those tasks the secretary can be triggered to do (i.e., that she does habitually). 'Necessary tasks' are those tasks that the head nurse believes are strategies, that is, those tasks she believes others usually do and are expected. Being able to cover those tasks then means that the head nurse has competences that are required to do those tasks. One can thus interpret this property as 'A head nurse has the required competences to do the activities that are strategies for her and that can be triggered for the secretary.' We can query which agents have the required competences for such activities.

```
SELECT ?agent ?sharedActivity ?competence
WHERE {
    a:Secretary a:belief ?personalActivity.
    ?personalActivity a:trigger ?contextElement.
    a:Secretary a:belief ?sharedActivity.
    ?sharedActivity a:requiredCompetence ?competence.
    a:Secretary a:hasCompetence ?competence.
    ?agent a:belief ?sharedActivity.
    ?agent a:hasCompetence ?competence.}
```

The formal reasoner returns the overall head nurse, but not the team head nurse. The property is thus false in the sense that not all the head nurses can cover for the secretary. The secretary usually contacts the ER, but although the team head nurse believes this is necessary she does not have the competence to do it herself.

Property 3 We make the property more precise by specifying 'needing equipment' and 'can help each other out'. Here 'need' means that one requires certain resources to do the actions they usually do. One way to 'help each other out' can

be to usually do actions that do not use the same resource, so that it is free to use for the other. One can thus interpret this property as ‘The staff usually does actions that are afforded by different resources than the actions they believe others usually do.’ To specify this into a query, we query if someone actually does use a resource that is needed by others. If the query returns nothing, we know our original property is satisfied.

```
SELECT ?staffmember ?personalActivityIDo
      ?sharedActivityOthersDo ?resource
WHERE {
  ?staffmember a:belief ?personalActivityIDo.
  ?staffmember a:belief ?sharedViewOnActivityIDo
  ?staffmember a:belief ?personalActivityOthersDo
  ?staffmember a:belief ?sharedActivityOthersDo

  ?personalActivityOthersDo a:strategy ?parentActivity.
  ?sharedActivityOthersDo a:affordance ?resource.
  ?personalActivityIDo a:trigger ?something.
  ?sharedActivityIDo a:affordance ?resource.
FILTER (?staffmember != a:Patient)
FILTER (?myActivity != ?othersActivity)}
```

The first part of the query sets up two activities on which an agent has a personal and shared view. The second part specifies that one of those activities is a strategy (i.e., something others usually do) and the other one a trigger (i.e., something the staff member usually does). It then queries if a resource is used for both of these activities. The third part specifies that staff members are not patients and that the activities should be different. When the activities are the same the agents are namely cooperating using the same resource to do the same activity. The query indeed returns nothing, showing that our staff members help each other out using their knowledge of what equipment the others need.

Property 4 and 5 These properties can not be expressed in our formalism. Property 4 has a deontic flavour, that is, what one should do. Our formalism express strategies, that is, what one usually does. It does not capture that there is a consequence of not adhering to the norm. In future work, we could research if the gained expressibility by a deontic extension of our formalism is worth the extra complexity. Property 5 mentions the length of stay, which is a common key performance indicator in operational research literature on optimizing patient flows in hospital departments. To express anything about how long the length of stay is we would need to formalize the dynamic parts of social practices. However our formalization is limited to reasoning about the static structure and so we are not able to express this property in the current terminology.

6 Related Work

Our work is related to other ontologies (or high-level descriptions) of agents that aim to capture sociality for social analysis. This section aims to explain that our

social practice formalization is unique in its richness of expressing the social world. It expresses the social world in a unique way, namely in terms of shared action-associations: associations with competences, values, context elements or other actions. Consumat is a meta-model for consumer agents [11]. The social is captured in that an agent can see the behaviour of other agents and choose to adopt it. In other words, agents share which actions they can see, but not associations they might make with those actions. Consumat mainly focusses on what comprises the agent instead of what comprises an action. The fact that our formalization expresses the shared associations of an action is what allows us to express the social properties described in Section 3.3. OperA is a framework for agent organizations [6]. The social is captured in shared ‘contracts’. Contracts consists of agents, roles, clauses and objectives. Agents can thus reason about the social world in these terms of these top-down prescriptions: the objectives and responsibilities of agents, instead of the bottom-up associations one makes with actions. OperA can better express deontic statements about what other should do, but cannot express the shared action hierarchy, affordances or required competence that is needed to express the aforementioned properties about the equipment others need or what others value. MAIA is a meta-model to capture agent institutions [8]. It builds on the assumption that, “while understanding and explaining individual behaviour is extremely complex, social rules or institutions are more elicitable”. It places the social in these shared rules and institutions. Our social practice ontology also considers some of these rules (strategies), but relate them to one practice instead of to one model. That is, each practice has a different set of relevant strategies. MAIA’s focus on institutions allows it to express deontic statements, but it can also not express the shared associations humans make with actions that allow one to express the properties about equipments others might need or what others value.

7 Conclusion and Future Work

This paper aimed to formalize the social dimension of an ER to be able to automatic identify social bottlenecks. Section 3 presented this social dimension and Section 4 showed that we could capture it in a precise, unambiguous ontology based on social practices. We used the Protégé tool and a formal reasoner to ensure the coherency of the ontology. This check means that the meta-model is satisfiable (there is a possible instance), the model is consistent (e.g., one does not claim two individuals are different and the same) and that the model is an instance of the meta-model. The OWL ontology with the meta-model and model is available online. It provides a basis for decision support tools for ER. For example, the tool could be used to identify activities where the staff does not show understanding of certain needs of the patient. Management could then educate the staff about this deficit and improve the wellbeing of both staff and patient. Our work can also be used as an example for formalizing the social dimension of other multi-agent systems. The formalization enables us to use a formal reasoner to keep track of a complex domain; something hard to do

analytically. In addition, it allowed us to gain precise insights about the formal relation between different concepts (as discussed in Section 4).

This paper focussed on giving a proof of concept of how to use a formal reasoner to verify a number of properties one could desire of an ER. Future work, could do a more extensive empirical study to support concrete claims about Herlev Hospital in Denmark or other hospitals. Although, the ontology is limited in expressing normative and dynamic statements, we showed that the ontology is unique in its ability to verify complex social statements about helping colleagues, the needs of patients and understanding what is important. An important aspect of real social systems is that actors can deviate from established practices. Such deviations are represented in social practices, which can initially represent established practices and then evolve over time as agents enact them. The social practice model that we have shown in this paper is static though and rather represent a snapshot of a social practice. In future work, we aim to extend the model to support evolution and expressing properties that have a time component. This paper demonstrates that the current ontology can already be used to ensure staff and patient wellbeing by identifying possible social problems.

Acknowledgements

This work is supported by the PhD project between PDC A/S and Technical University of Denmark. We thank the institute Region H, which manages the hospitals in the Danish capital region, for providing input for the use case. We would also like to thank Helle Hvid Hansen, Jørgen Villadsen and Jordi Bieger for comments on the ideas described in this paper and comments on a draft.

References

1. Allen, D.: Re-reading nursing and re-writing practice: Towards an empirically based reformulation of the nursing mandate. *Nursing Inquiry* **11**(4), 271–283 (2004). <https://doi.org/10.1111/j.1440-1800.2004.00234.x>
2. Berardi, D., Calvanese, D., Giacomo, G.D.: Reasoning on UML Class Diagrams using Description Logic Based Systems. In: Proceedings of the KI-2001 Workshop on Applications of Description Logics (KIDLWS'01). pp. 1–12 (2001). <https://doi.org/10.1.1.177.1920>
3. Berg, M., Aarts, J., Van der Lei, J.: ICT in Health Care: Sociotechnical Approaches. *Methods of Information in Medicine* **42**(4), 297–301 (2003). <https://doi.org/10.1267/METH03040297>
4. Bock, C., Fokoue, A., Haase, P., Hoekstra, R., Horrocks, I., Ruttenberg, A., Sattler, U., Smith, M.: OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition) (12 2012), <http://www.w3.org/TR/2012/REC-owl2-syntax-20121211/>
5. Crawford, S.E.S., Ostrom, E.: A Grammar of Institutions. *Political Science* **89**(3), 582–600 (2007). <https://doi.org/10.2307/2082975>, <http://www.jstor.org/stable/2082975>
6. Dignum, V.: A Model for Organizational Interaction. Ph.D. thesis, SIKS (2004)

7. Dignum, V., Dignum, F.: Contextualized Planning Using Social Practices. In: Ghose, A., Oren, N., Telang, P., Thangarajah, J. (eds.) *Coordination, Organizations, Institutions, and Norms in Agent Systems*, pp. 36–52. Springer International Publishing, Cham (2015)
8. Ghorbani, A., Bots, P., Dignum, V., Dijkema, G.: MAIA: A framework for developing agent-based social simulations. *Jasss* **16**(2), 1–15 (2013). <https://doi.org/10.18564/jasss.2166>
9. Harris, S., Seaborne, A.: SPARQL 1.1 Query Language (2013), <https://www.w3.org/TR/2013/REC-sparql11-query-20130321/>
10. Horrocks, I., Patel-schneider, P.F., Boley, H., Tabet, S., Grosz, B., Dean, M.: SWRL : A Semantic Web Rule Language Combining OWL and RuleML. W3C Member submission (May 2004), 1–20 (2004)
11. Jager, W., Janssen, M.: An updated conceptual framework for integrated modeling of human decision making: The Consumat II. *Eccs* 2012 p. 10 (2012)
12. Mans, R.S., Aalst, W.M.P.V.D., Vanwersch, R.J.B.: *Process Mining in Healthcare Evaluating and Exploiting Operational Healthcare Processes* (2015). <https://doi.org/10.1007/978-3-319-16071-9>, http://alexandria.tue.nl/extra1/afstvers1/wsk-i/Steeg_2015.pdf%5Cnhttp://link.springer.com/10.1007/978-3-319-16071-9
13. Norling, E.: Don't Lose Sight of the Forest: Why the Big Picture of Social Intelligence is Essential. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. pp. 984–987. No. May (2016)
14. Poel, I.v.d., Royackers, L.: *Ethics, Technology, and Engineering: An Introduction*. John Wiley & Sons (2011), <http://eu.wiley.com/WileyCDA/WileyTitle/productCd-EHEP002302.html>
15. Reckwitz, A.: Toward a theory of social practices: A development in culturalist theorizing. *European Journal of Social Theory* **5**(2), 243–263 (2002). <https://doi.org/10.1177/1368431022225432>
16. Rojas, E., Munoz-Gama, J., Sepúlveda, M., Capurro, D.: Process mining in healthcare: A literature review. *Journal of Biomedical Informatics* **61**, 224–236 (2016). <https://doi.org/10.1016/j.jbi.2016.04.007>, <http://dx.doi.org/10.1016/j.jbi.2016.04.007>
17. Rumbaugh, J., Jacobson, I., Booch, G.: *The unified modeling language reference manual*. Addison-Wesley (2005), <https://dl.acm.org/citation.cfm?id=993859>
18. Schatzki, T.R.: *Social Practices. A Wittgensteinian approach to human activity and the social*. Cambridge University Press (1996)
19. Shove, E., Pantzar, M., Watson, M.: *The Dynamics of Social Practice: Everyday Life and How it Changes*. SAGE Publications, London (2012). <https://doi.org/10.4135/9781446250655.n1>
20. Sørup, C.M.: *Development of a Generic Performance Measurement Model in an Emergency Department*. Ph.D. thesis, Technical University of Denmark (2015)
21. W3C OWL Working Group: *OWL 2 Web Ontology Language Document Overview (Second Edition)* (12 2012), <http://www.w3.org/TR/2012/REC-owl2-overview-20121211/>
22. Waring, J.J., Bishop, S.: Lean healthcare: Rhetoric, ritual and resistance. *Social Science and Medicine* **71**(7), 1332–1340 (2010). <https://doi.org/10.1016/j.socscimed.2010.06.028>, <http://dx.doi.org/10.1016/j.socscimed.2010.06.028>
23. Wood, W., Neal, D.T.: A new look at habits and the habit-goal interface. *Psychological review* **114**(4), 843–863 (2007). <https://doi.org/10.1037/0033-295X.114.4.843>