Who's That? - Social Situation Awareness for Behaviour Support Agents
A Feasibility Study

Ilir Kola[1(✉)], Catholijn M. Jonker[1,2], and M. Birna van Riemsdijk[3]

[1] Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands
{i.kola,c.m.jonker}@tudelft.nl
[2] Leiden Institute of Advanced Computer Science, Leiden, The Netherlands
[3] Human Media Interaction Lab, University of Twente, Enschede, The Netherlands
m.b.vanriemsdijk@utwente.nl

Abstract. Behaviour support agents need to be aware of the social environment of the user in order to be able to provide comprehensive support. However, this is a feature that is currently lacking in existing systems. To tackle it, first of all we explore literature from social sciences in order to find which elements of the social environment need to be represented. We structure this knowledge as a two-level ontology that models social situations. We formalize the elements that are needed to model social situations, which consist of different types of meetings between two people. We conduct an experiment to evaluate the lower level of the ontology using feedback from the subjects, and to test whether we can use the data to reason about the priority of different situations. Subjects found our proposed features of social relationships to be understandable and representative. Furthermore, we show these features can be combined in a decision tree to predict the priority of social situations.

Keywords: Socially aware agents · Social situation modelling · Knowledge representation

1 Introduction

Artificial agents that support people in their daily lives, for example to live healthier lifestyles or help them in the execution of daily tasks, are becoming a reality (e.g. [32,43]). Such behaviour support agents need to be aware of a user's *social context* to function effectively [46]: a user's social network may need to play a role in providing support, and a user's activities may involve other people which affects the type of support that is needed [41]. For instance, an app that helps its user be more punctual might send reminders at different intervals when it sees that a meeting is approaching. However, not all meetings have the same

This work is part of the research programme *CoreSAEP* (project no. 639.022.416), supported by the Netherlands Organisation for Scientific Research (NWO).

© The Author(s) 2020
L. A. Dennis et al. (Eds.): EMAS 2019, LNAI 12058, pp. 127–151, 2020.
https://doi.org/10.1007/978-3-030-51417-4_7

priority: for most people, being on time for a job interview is more important than being on time for an informal dinner with friends. Effective support may require taking this into account in the frequency or type of reminders that are generated.

Existing behaviour support agents however mostly focus on modelling internal aspects of the users (e.g. their goals, values, abilities, etc.) [36,44], while paying less attention to users' social context. In this paper we take first steps towards developing a *generic framework* that enables behaviour support agents to take into account the user's social environment in order to provide personalized and socially-aware behaviour support [46].

The main idea underlying our approach is to take research on situation awareness, which offers ways to model and reason about the physical environment, and adapt it for realizing *social situation awareness*. Specifically, we take the well-known situation awareness model by Endsley [17] as a starting point. Endsley's model distinguishes three levels of situation awareness: 1) *perception* of relevant elements in the environment, 2) *comprehension* to understand their significance, and 3) *projection* towards future states of the environment. Inspired by these levels, we put forward the idea that a behaviour support agent should similarly be able to represent relevant aspects of social situations, be able to reason about their meaning, and lastly project how these situations will affect the behaviour of the user. These three levels are in line with the classic sense-reason-act cycle in multi-agent systems.

While there are many socially relevant dimensions to behaviour support, in this paper we focus on handling *social settings* such as meetings or social gatherings. Moreover, we focus on behaviour relevant for arranging these social settings, rather than how to behave whilst participating in one. One may think of a personal assistant agent that can schedule social events for its user [41], or an agent to support people with cognitive impairments in arranging their social life. Furthermore, we need to determine which dimensions of a social situation may be used to interpret their meaning, i.e., what is the "output" of the comprehension process. In this case we focus on *priority* of social situations. We expect that priority, among other things, may be used for dealing with conflicts in a user's schedule. Putting this together, in this paper we address the following research questions and hypothesis, corresponding with the three levels of situation awareness:

- **RQ1: Perception -** Which features can be used to describe a social situation from the perspective of a user for the purpose of behaviour support?
- **RQ2: Comprehension -** How can features of a social situation be used to assess its priority?
- **H3: Projection -** Priority of social situations can be used for resolving conflicts between two social settings if they cannot both be attended.

While these research questions and hypothesis guide the work presented in this paper, we do not aim to provide definitive answers here. Rather, as this is a novel research direction, our aim is to assess the feasibility of the approach as a basis for future work that considers other dimensions besides priority, as well as a more extensive investigation into their translation to support actions by the agent.

Addressing these questions involves creating knowledge structures and reasoning techniques for representation and interpretation of social situations, as well as evaluation with users. We further detail this approach and the envisaged software architecture for our support agent in Sect. 2. We present a knowledge structure for describing features of social situations in Sect. 3. We present our user study to evaluate this knowledge structure and gather data for addressing RQ2 and H3 in Sect. 4. Our reasoning model for addressing RQ2 is presented in Sect. 5. We conclude the paper and discuss our findings in Sect. 6.

2 Research Approach and Agent Architecture

The overall objective of this work is to assess the feasibility of realizing social situation awareness for behaviour support agents based on the three levels of situation awareness of Endsley [17]. For this reason, we touch on each of these three levels in this work (albeit less comprehensively for the higher levels), i.e., we take a "breadth-first" approach, rather than first going into depth on the first level. In this way we get a sense of how the different levels of the framework could work together to achieve social situation awareness early on in the research, and it allows us to identify aspects that require a more in-depth study in follow-up research. Specifically, we address the research questions and hypothesis in the following way:

- RQ1: Which features can be used to describe a social situation from the perspective of a user for the purpose of behaviour support?
  - Model building: based on research in social sciences, we propose an ontology for modelling the high-level structure of social situations, as well as a set of low level features that can be used to describe daily life social situations (Sect. 3).
  - Evaluation: Assessing whether the social features identified in the modelling step are suitable, consists of two parts: i) assessing the understandability and expressivity of these features for users; this is important, since we envisage that we will (partly) elicit these features from users through interaction with the support agent, and explaining the support agent's actions to the user requires that these features are meaningful to users (Sect. 4); ii) assessing the usefulness of these features for situation comprehension; this is assessed via RQ2 (Sect. 5).
- RQ2: How can features of a social situation be used to assess its priority?
  - Model building: One may envisage different ways of building a model that can take features of a social situation and derive a corresponding priority, for example by pre-specified rules, through machine learning, or a combination. Since an important requirement for this model is its explainability for users, in this paper we choose a learning method that

yields an interpretable model: decision trees. To create this decision tree, we collect data from people via a user study (Sect. 4), and then use the data to learn a decision tree that predicts priority of social situations (Sect. 5).
    - Evaluation: We evaluate the predictive capacity of the decision tree by taking a test data set from the data collected for building the model, and evaluating its capacity in predicting the right priority for a specific event based on information about social features of the situation.
- H3: Priority of social situations can be used for resolving conflicts between two social settings if they cannot both be attended.
    - Data collection: First, we ask subjects about their relationship with people in their social circle. Then, we present them with social situations involving these people, and ask them what priority they would assign to these situations. Lastly, we show them pairs of these situations and ask them which one they would attend if the meeting times would overlap and they had to choose only one (Sect. 4).
    - Hypothesis testing: To test this hypothesis, we check whether the proportion of meetings with higher priority that was chosen when breaking the ties is higher than chance.

Figure 1[1] depicts a high level architecture of how our proposed behaviour support agent can be used in practice. The first part of the work consists in learning a model which given data from different social situations (in our case, the experiment data), it learns priority rules based on the answers of the participants. When the user is faced with a future social situation, it gives the behaviour support agent a description of the situation (*situation cues*) and relationship with the other person (*social background features*). The agent uses this information, as well as the learned priority rules, to reason about the priority of this situation. In future work, the priority level will be fed to a support reasoner, which will then output a support action to be of assistance to the user. In this work, we hypothesize that priority can be used to break ties when different meetings overlap. In that case, the support reasoner can compare the priority of the different meetings, and suggest to the user which one to attend.

3 Modelling Social Situations

In this section we outline which features can be used to describe a social situation from the perspective of a user of the behaviour support agent. We distinguish between a description of the main components, i.e., the overall structure of a social situation (Sect. 3.1) which we refer to as the upper ontology following [24], while the concrete features of the social situation that are the result of the perception process are described in a lower ontology (Sect. 3.2).

[1] Icons used in the architecture were made by Freepik and retrieved from www.flaticon.com.

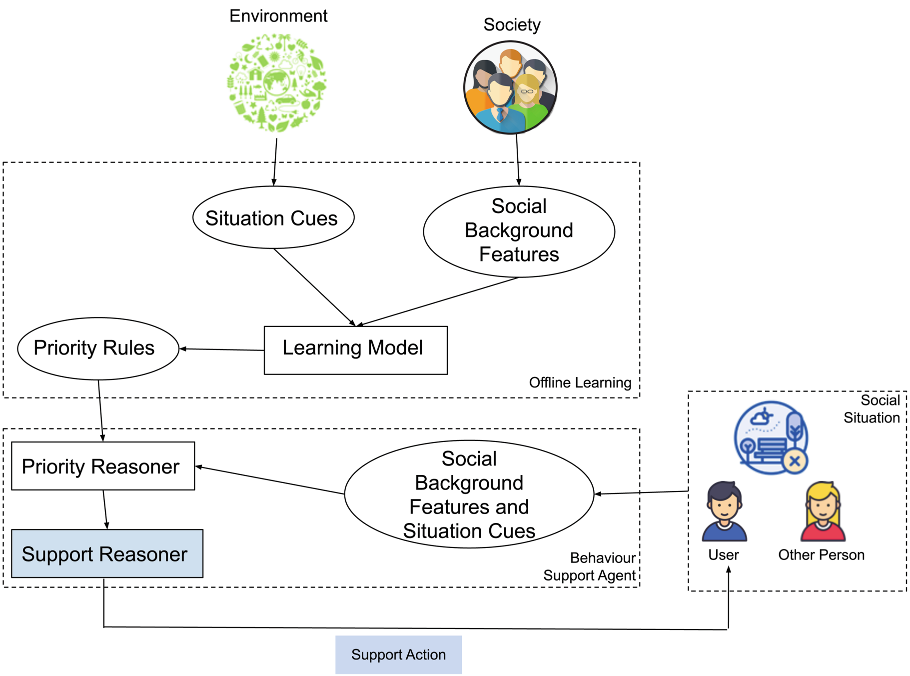

Fig. 1. High level architecture of the proposed approach. Boxes marked in blue are parts which we do not explicitly tackle in this paper. (Color figure online)

3.1 Structure of Social Situations: Upper Ontology

Research in social psychology by Rauthmann and colleagues [37] proposes that features of situations can be discussed on three different levels: *cues*, which are physical and objective elements (who is present, what activity is taking place, etc.), *psychological characteristics*, which are dimensions that can be used to describe situations (such as duty, intellect, etc.), as well as *classes*, which are abstract types of situations (such as social situations, work situations, etc.). For the scope of this work, we will focus on situation cues and classes, since these are concrete concepts that can be elicited from the user, i.e., that are the result of the perception process. Psychological characteristics, and how to automatically infer them, will be explored in future work.

Cues in turn can be divided into three categories according to [37]: persons, events/activities, and locations. Saucier et al. [39] identify similar categories in an experiment in which students describe their daily situations, namely locations, associations (i.e. people/interactions), as well as actions and positions. Thus we can see that in the literature information about people in the situation is considered to be a specific kind of situation cue. Since in this paper we focus on modelling *social* situations, meaning that the relation to the people in the situation is of specific interest, we decide to model people separately from other situation cues. This is in line with other work in the field of socially intelligent technologies [1,2].

Cues. The literature identifies essentially two remaining types of cues [37,39], when we separate information about people from other situation cues: *location* and *activity*. In this paper we also model the situation class as a type of cue, which we refer to as the *setting*. Furthermore, we introduce a number of additional cues that we consider specifically relevant for comprehension of organized events, as we focus on in this paper. In particular, we represent the *frequency* with which an event takes place. This variable is not explicitly mentioned as a situation cue in the literature, however some situation taxonomies, e.g., [33], suggest typicality as one of the psychological characteristics of the situation. Moreover, we represent the *time* at which the event takes place, as well as the *initiator* of the event, since we expect this may influence the priority of the meeting.

People. For reasons of simplicity, in this work we focus on dyadic social relationships, i.e., we concern ourselves with social situations involving two people. In our case, one of the people will be the user of the behaviour support agent. This means that the information about the social relation is modelled *from the perspective of the user*.

We model the social relationship by identifying a set of features that characterize this relationship. We distinguish between *social background features* and *situation-specific social features*. The former concern features that describe aspects of the relationship in general, while the latter describe aspects that are specific to the situation at hand. We distinguish two kinds of social background features, namely *structural features* and *personal features*. The former concern what may be referred to as "objective" characteristics such as the user's role in relation to the other person, while the latter concern "subjective" relationship characteristics from the perspective of the user, such as the quality of the relationship. This distinction is in line with research in social science on relationships in organizations [30] and social support [25], which considers the difference between relationship characteristics that are derived from formal requirements of a role, and interpersonal characteristics. These features are further detailed in Sect. 3.2.

Putting this all together, Fig. 2 offers a schematic representation of the upper ontology.

Related Work. Context and situations are well studied concepts in computer science. Kokar and colleagues [27] present an ontology for formalization of situations based on the situation theory developed by Barwise [5] and extended by Devlin [10]. This formalization is compatible with the interpretation of situation awareness provided by Endsley [18], which also forms the basis of our work. Yau and Liu [48] offer another ontological approach that models situations for pervasive computing applications. They differentiate between situations, defined as

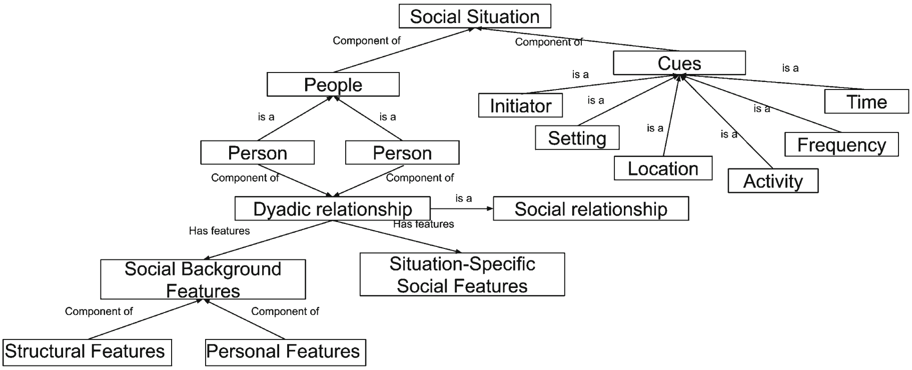

Fig. 2. Schematic representation of an upper ontology of dyadic social situations.

"a set of contexts in the application over a period of time that affects future system behavior" and contexts, defined as "any instantaneous, detectable, and relevant property of the environment, system, or users". Their ontology is based on this division, and they specify a context layer, which models context definition and contextual data, and a situation layer which is built on top of the context layer and aggregates context into situations. This forms the core of their upper ontology, whereas the elements of the lower ontology can be specified depending on the domain. Their definition of context can be compared to our notion of situation cues. However, these approaches are very abstract in the concepts used in the ontology since they focus on modelling a generic type of situations. Building the lower level ontologies, specifically concerning the modelling of social situations as we focus on in this paper, is not a trivial task.

Zavala and colleagues [50] offer a framework which can be used to build *place-aware mobile applications*. To do so, they build a place ontology which models the concept of place not only as a geographical location, but also in terms of activities that occur there. For instance, someone can have an office in two different cities, but both of them would count as a *workplace* since similar activities occur there. This is comparable to the cues "location" and "setting" in our ontology. In Murukannaiah et al. [31] this approach is extended and social circles are learned based on the places in which people are meeting: following the previous example, people meeting in workplaces would be classified as colleagues. This can be viewed as a kind of structural relationship feature, as we refer to it in our ontology. Similar to our work, their approach goes beyond modelling very abstract concepts for representing generic situations. However, the concept of places and associated types of relationship is just one aspect relevant to comprehending social situations. Our approach aims at providing a more comprehensive knowledge structure for modelling social situations, as well as development of methods for interpreting these.

Another related line of research is work on modelling and reasoning about *social practices* [11,14]. In [14], social practices are represented by distinguishing physical context (resources, places, actors), social context (social interpretation, roles, norms), activities, plan patterns, meaning and competences. Physical context and activities are comparable to what we refer to as cues of a social situation, while social context in our case concerns the modelling of people. Meaning can be compared to our second level of situation awareness, i.e., comprehension. Thus while the type of notions we use for modelling social situations are broadly comparable to what is used in research on social practices, our starting point is different. In social practice modelling, the starting point is the social setting, e.g., a classroom [11], for which the norms and expected activities are explicitly modelled *independent of the participating agents*. Then deliberation techniques are needed to allow agents to determine how to achieve their goals, taking into account the (given) norms of this social setting [14]. In our work the starting point is the social relation between the (human) *agents*. For this reason we go in detail regarding the modelling of social background features (Sect. 3.2) that characterize from the (subjective) user's point of view their relation with the other person in the social situation. Based on these features, we then interpret in a bottom-up way the social situation in terms of more abstract general characteristics, in this case priority of a social event. From that we then determine appropriate support actions for the user. Moreover, since our aim is to create behaviour support agents for *people*, we develop our models taking into account results from user studies.

In our previous work [28] we provide an extension of the ontology of [27] with relations that support modelling social relationships, and explore how these can be used for decision making in social situations. However, in that paper we model social relations based on only four abstract relationship types from [19] that can be used to model social decision making: communal sharing, authority ranking, equality matching, and market pricing. These can be viewed as a type of structural relationship feature. However, these do not capture personal features that describe more subjective aspects of interpersonal relationships. Moreover, in that paper we do not investigate comprehension of a social situation based on these features, but rather model decision making directly using pre-specified rules.

3.2 Features of Social Situations: Lower Ontology

In this section we go more in detail regarding the modelling of situation cues, and we introduce features of dyadic social relationships that a behaviour support agent can use to model daily life social situations of a user. The list of features presented in this section is not exhaustive, and depending on the type of behaviour support different features may be relevant. However it highlights the type of features that may be considered, and serves as an example of the concrete features that can be used. Moreover, we use these features to model the scenarios in our experiment.

We represent features of social situations by means of relations over situation instances ($\mathcal{SI}$) and dyadic social relationships ($\mathcal{A} \times \mathcal{A}$ where $\mathcal{A}$ is the set of people) for cues and social features respectively, and a domain ($\mathcal{D}$) that specifies the value-ranges the feature can take. This is in line with situation theory ontology [27] in which the modelling of perceived aspects of a situation is done by means of so-called *infons* which describe the relations between objects in a situation. The appropriateness of the chosen value-ranges is also subject to evaluation, and may be changed depending on the domain.

Cues. For simplicity in this paper we focus on three out of six cues that have been introduced in Sect. 3.1: the initiator of an event, the setting of a social situation, and frequency of the event. A good starting point for modelling locations and activities can be the work of Zavala et al. [50].

The initiator is a person from the set $\mathcal{A}$, or `none` if no initiator is identified. For the selection of types of setting of a situation we choose common situation classes that users may face in their daily life. In this paper, we base the types of settings on Pervin [35], who identifies work situations, family situations, friends/recreation situations, and private recreation situations. We omit the latter since we are concerned with social situations, and add sports activity as a specific type of setting. The situation classes proposed in Rauthmann et al. [37] can also be clustered into these settings. We distinguish two frequencies, `regular` and `occasional`. While more fine-grained distinctions can be made, we expect that this broad categorization suffices in many cases. We list the corresponding relations in Table 1 below.

Table 1. Relations to model cues of social situations. For a relation $\langle name \rangle$, set of situation instances $\mathcal{SI}$ and domain $\mathcal{D}$, the relation is defined as $\langle name \rangle : \mathcal{SI} \times \mathcal{D}_{\langle name \rangle}$.

Relation name	Domain ($\mathcal{D}$)
event_initiator	$\mathcal{A} \cup$ {`none`}
setting	{`work_related, casual_meeting, sports_activity, family_related`}
event_frequency	{`regular, occasional`}

Social Background Features. While there is a lot work in the social sciences on understanding social relationships, in this paper we mainly use the following two lines of work as the basis for selecting structural and personal social features for our model. First, Kahn and Antonucci [3,4,25] explore the role of social relations as a form of social support for (elderly) people. Enabling social support is an important purpose of the behaviour support agents we aim to create [46]. We select our structural features mainly from this line of work. Second, social relations are also considered from the organizational point of view. Specifically, we use the work of Mainela [30] which gives an overview of types and functions of social relationships that can be relevant in the organization of a joint venture. Organizational relationships are an important type of relation that our behaviour support agent may take into consideration. We select our personal features mainly from this work.

Structural Features. Kahn and Antonucci conceptualize support systems as a so-called Convoy model - three concentric circles representing three levels of closeness between the supported person and their "convoy" of supporters. Different aspects of the relationship are considered in order to establish someone's position within the convoy model. The Convoy model [4] distinguishes between structural (age, sex, years known, proximity, contact frequency, relationship (role)) and functional characteristics (types of support received and provided) of social support networks.

For this paper we use *role*, *contact frequency*, and *default geographical distance* (proximity) as structural features. The feature *role* refers to the role of the other person towards the user in dyadic relations. Knowing this is important since it can help inferring the expectations that come with the role. The range of roles we use is taken from the general social survey [8]. The geographical distance refers to the physical proximity of the two actors in terms of their default home location. Proximity can influence the relationship of two people since it affects how often they can see each other. For the range we opted to measure distance in terms of time that it usually takes to get to that person.

Besides the above three structural features, we introduce a fourth one, namely *hierarchy*, to express the type of relation between the user and the other person. Hierarchy affects the power dynamics between the first and second actor. Higher (respectively same and lower) means that the other person is higher up (resp. at the same level, and lower) in the hierarchy than the user. In case there is no hierarchy amongst the actors, this is indicated by "n.a.". We expect this feature to be relevant when assessing the priority of meetings, especially for users who are in working relations, or actors that come from a culture with some sort of caste system. More information on the concept of hierarchical ranking can be found in, e.g., [19,47].

Personal Features. The first of our personal features is also taken from the Convoy model [3]. In addition to structural and functional aspects of relationships, this paper emphasizes the importance of *relationship quality* in characterizing social relations. The remaining three personal features we consider in this paper are taken from Mainela [30]. The paper gives an overview of how types of social relationships in business dyads have been characterized in the literature. For example, Granovetter [22,23] talks about strong ties and weak ties in work relationships. The strength of a tie in a network depends on aspects such as the amount of time spent on it, the emotional intensity, the intimacy, and the reciprocity. Furthermore, the author argues that ties are stronger when the level of acquaintance is deeper.

From the list of features for characterizing social relations identified through the literature study of Mainela, we select three, namely *acquaintance depth* [22] of the user towards the other person, *level of formality* of the relationship [38], and *trust* [45] of the user towards the other person as personal features. These features can inform the expectations of the relationship between user and the other person, and consequently are relevant for comprehending social situations.

Other features mentioned by Mainela can be used to distinguish different types of social relationships in a business context, but seem too specific for social situation awareness of our envisaged behaviour support agent, e.g., legal questions, attendant consequences, activation of a relation, outcome expectations, and scope of economic issues. The features continuity of interaction and amount of time spent are closely related to event frequency, contact frequency and acquaintance depth. Features like personal nature, intimacy and emotional intensity seem closely related to level of formality and acquaintance depth. Finally, reciprocity may also be relevant for our purposes, however refers more to functional aspects of the relationship and may be difficult to characterize directly in these terms by users. Therefore we leave it out in this paper.

We summarize these social background features in Table 2 below. The range of some features is $\texttt{Likert}_5$, which denotes a 5-point Likert-type scale, where 1 is the lowest/most negative value and 5 the highest/most positive value.

Table 2. Relations to model social background features of social situations. The upper part concerns structural features, the bottom part personal features. For a relation $\langle name \rangle$, and domain $\mathcal{D}$, the relation is defined as $\langle name \rangle : \mathcal{A} \times \mathcal{A} \times \mathcal{D}_{\langle name \rangle}$ where $\mathcal{A}$ denotes the set of persons.

Relation name	Domain ($\mathcal{D}$)
role	{partner, parent, sibling, child, extended_family, coworker, neighbor, friend, supervisor, group_member, other}
contact_frequency	$\texttt{Likert}_5$
def_geo_distance	{0-1 h, 1-2 h, 2-4 h, flight needed}
hierarchy	{higher, same, lower, n.a.}
rel_quality	$\texttt{Likert}_5$
acq_depth	$\texttt{Likert}_5$
rel_formality	$\texttt{Likert}_5$
trust	$\texttt{Likert}_5$

Situation-Specific Social Features. Several of the social background features may have a situation-specific variant, for example if you go to a basketball game with your boss, in that situation you are both team-mates, and if you are the captain you are the one holding a higher hierarchy level in that situation. However for reasons of simplicity we do not further elaborate on these in this paper.

We do introduce another situation-specific social feature, which we call the *help dynamic*. It refers to whether in the specific event the user is giving to or receiving help from the other person. The fact that they have to give or receive help can influence how obligated the actors feel to attend a certain event. It is defined as a relation $help\_dynam : \mathcal{SI} \times \mathcal{A} \times \mathcal{A} \times \mathcal{D}_{help\_dynam}$, where $\mathcal{D}_{help\_dynam} = \{\texttt{giving, receiving, neither}\}$.

Related Work. Different aspects of modelling social relationships have been studied in sub-fields of multi-agent systems. In particular, when talking about organizations of agents, "role" is one of the central concepts. In the OperA model [13], agents form societies with different organizational structures, and they take up roles in these societies. These roles, in combination with social contracts, define what an agent should and should not do. Singh [40] follows a similar approach, and proposes that "Org(anization)s are finely structured through the notion of a role, which codifies a set of related interactions that a member of an Org may enact". D'Inverno and colleagues [16], in their quest to weave a fabric for socially aware agents, also introduce the concept of roles in order to represent agents in the context of a social setting. Roles in these works are used to describe, design and understand interactions in an abstract and re-usable sense, independent from the agents that will eventually play the roles. In our case we combine abstract information about roles with information about the concrete relation between the user and the other person, i.e., between the specific (human) agents in the interaction, in order to assess how best to support the user this social situation.

The notion of hierarchy is used in [13] to describe a type of relation between roles in an organization. Although not the same thing, hierarchy can be connected to the notion of power. Pereira and colleagues [34] argue for the importance of modelling social power into the decision making of cognitive agents. The importance of modelling social power is also proposed in [12].

Another well studied concept within the multi-agent systems field is trust. Mostly, it is considered from the point of view of software agents trusting each other. The focus is on determining the level of trust in another agent by taking into consideration the agent's previous interactions with another agent, or by relying on other agents' opinions about that agent [20,49]. In our case, once we have information about the trust the user has towards the other person, we use it for interpreting the social situation and allowing our support agent to determine the appropriate support actions in this situation.

The virtual agents research area has also studied modelling and use of various features that describe social relationships. Zhao and colleagues [51] argue for the importance of representing rapport in a virtual agent that interacts with a human. Rapport is a feeling of connection and closeness to another person, which can be compared with depth of acquaintance. Dudzik and colleagues [15] provide a review of literature that deals with contextual features of human emotion perception for automatic affect recognition. As contextual factors they identify characteristics of the sender or receiver of the emotion, such as age, gender and occupation, as well as situation features such as cause of the emotion, conversation content and language, information about the conversation partner in the social interaction, location, and lighting conditions during the interaction. Our work is complementary in that it focuses on characterizing the social relationship itself between people in the social situation, and from that derive higher-level understanding of the social situation, in this case in terms of its priority.

Thus our framework for modelling social situations includes a number of features that have been studied in various parts of the agent systems literature. Based on social science literature we add several features that are specifically relevant for characterizing human social relations, such as contact frequency, geographical distance, and relationship formality. Moreover, our work differs from existing work in multi-agent systems in that we investigate how we can *combine* features of social situations for the purpose of comprehension in order to allow an agent to provide appropriate socially-aware support.

4 User Study

In order to evaluate how well we can use our proposed low level features to model and interpret daily social situations, we conducted a pilot experiment in which subjects had to answer a survey about the social relations in their life [29]. The survey consisted of three parts, through which we explore RQ1 and evaluate H3. Furthermore, we use the data from survey to create and evaluate a model that addresses RQ2. We present our experimental setting in Sect. 4.1 and our results in Sect. 4.2.

4.1 Experimental Setting

Pilot Subjects. We tested 20 subjects (15 male, 5 female) who answered to all three parts of the experiment. Subjects were university employees (mostly PhD candidates). The average age was 31.1 years old (SD = 7.6yo).

Design and Procedure.[2] The experiment was implemented as an online survey, and consisted of three parts. In *Part I – Perception*, subjects were asked to think about six people from their social circle. For the purpose of the study, they were instructed to select at least one family member, one friend, and one person who had a higher hierarchy level than them. In follow-up research, we will also ask for information on relationships with people lower in the hierarchy. For each of these people, subjects were asked to provide all social background features (Sect. 3.2). The first part was concluded with an evaluation section in which the subjects were asked whether the questions were understandable, whether the amount of questions was appropriate, and how well they thought the questions represent their social relationship with someone. Through these questions, we test how understandable and expressive our proposed features are (RQ1). Furthermore, they had the option to propose more aspects of social relationships which they thought are relevant.

[2] The questions for each part of the experiment can be found in the Appendix.

In *Part II – Comprehension*, subjects were shown 20 scenarios of daily life social situations. Each scenario involved one of the six people that subjects had mentioned in Part I, selected randomly[3]. We made the study subject-specific to enable them to reflect on their own relationships, instead of presenting them with hypothetical relationships. Scenarios consisted of different parameters of the situation cues and situation specific features of social relationships. A scenario could represent a social situation such as:

> "You have invited *Person X* for a work meeting on Tuesday morning because you need some feedback on your recent project".

In this case it is a work setting, the event is occasional, the subject is the initiator and he/she is expected to receive help. For each scenario, subjects were asked about the priority of the meeting, how obligated they would feel to attend the meeting and how much they would enjoy it. We need the information on priority to answer RQ2. Obligation and enjoyment were asked for exploratory purposes to inform future research. Furthermore, subjects were asked how they think the other person would answer these questions. This was done because in future work, we want to explore the reciprocity of these decisions. Lastly, they were asked about the likelihood of that scenario happening in their daily life in order to assess the appropriateness of the scenarios we have chosen. Subjects had to answer on a 5-point Likert scale. In order to assess priority, they were instructed to take into account how difficult it would be for them to cancel the meeting, how important they think it is to be punctual, and any other thing they would consider relevant.

In *Part III – Projection*, scenarios were paired randomly and subjects were asked which of the two meetings would they choose to attend in case of a conflict between the two scenarios meaning that they could not attend both meetings. We will use this information to evaluate H3. Furthermore, they were asked what reason would they give to the person whose meeting they were canceling: the real reason, some other reason, or no reason. This was asked in order to have some more insight in case our hypothesis is not corroborated from the data. Each subject was presented with six pairs of scenarios.

4.2 Results

In this subsection, we will present and discuss the results of each part of the experiment separately.

Part I – Perception. The selected people from the subjects' social circle had an average age of 37.6 years old (SD = 13.55 yo). They were mostly friends (29%), followed by people from work (18% supervisors and 10% coworkers) and family

[3] Apart from the scenarios in which a family setting or a higher hierarchy work setting were being tested, which were restricted to family members and people with higher hierarchy, respectively.

members (11% parents, 8% siblings and 7% members of the extended family). Partners consisted of 10% of the selected people. Overall 74% of the people were not in a hierarchical relation with the subjects, 22% were on a higher level and 4% on a lower level. 36% lived within an hour of distance from the subjects, 18% between 1–2 h, 4% between 2–4 h, and for the remaining 32%, the subjects would need to take a flight in order to meet them. The subjects' answers for social background features that have a Likert-scale as the domain are shown in Table 3 below.

Table 3. Percentage of subjects that gave each specific answer for different social background features. The answer options were Likert-type scale values ranging from 1 to 5. For relationship quality 1 = very negative and 5 = very positive. For the rest, 1 = very low and 5 = very high.

Feature\Answer	1	2	3	4	5
Contact frequency	0	23.82	25.59	25.59	25
Relationship quality	2.06	5	10.88	49.41	32.65
Acquaintance depth	0	15.59	34.12	25.59	24.71
Relationship formality	46.76	16.18	26.18	7.06	3.82
Trust	1.76	1.47	25	37.25	34.41

As seen in Table 3, subjects mostly choose people with whom they have strongly positive relationships. Furthermore, they chose people whom they trust, and the relationships have a low level of formality. In future work, in order to have more representative data from a larger variety of relationships, we will control some features when asking the subjects to think of people from their social circle. For instance, we will ask some subjects to think about a coworker with whom they do not have a positive relationship.

The evaluation questions (all posed with a 5-point Likert scale in possible answers) showed that the subjects found the questionnaire understandable, with an average of 4.59 (SD = 0.51). The number of asked questions was appropriate (the average answer was 3, SD = 0.61, on a 5-point scale where 3 = `appropriate`). When asked how much this information represents their relationship with someone (Likert range from 1 = `very little` to 5 = `very much`), the average answer was 3 (SD = 0.79), confirming that social relationships have subtle aspects not captured in our questionnaire. Whether we need to add more features, depends on the strength of the correlations between the current features and the choices the subjects make in Part II of the questionnaire. The subjects (mostly being PhD students), seemed to understand this point, as some subjects indicated that the answer to this question depends on the purpose of the study. This is something that we will take into account in future experiments.

When asked whether they could think of additional aspects of social relationships which should be present in the survey, 35% of subjects answered with

"Yes". Some of the suggestions included: dependability, understanding, fun, respect, how important is the other person, common interests, etc. However, none of the suggestions appeared consistently.

Part II – Comprehension. In this section subjects were asked to evaluate different scenarios with respect to their priority (and additionally obligation and enjoyment). Subjects mostly give a high priority to the meetings, with 37% of scenarios being assigned a 5, 41% a 4 and 16% a 3, with only 6% having a 1 or a 2. This was expected given that scenarios included people with whom the subjects have a close and positive relationship. This is also reflected in how much they enjoy these meetings (65% of scenarios being assigned a 4 or a 5). For obligation, the results were more balanced, with 14% of scenarios being assigned a 2, 21% a 3, 37% a 4 and 25% a 5. The average likelihood of the scenarios was 3.14 ($SD = 1.42$), which means the scenarios were relatively likely despite being chosen randomly in terms of the combination of person with whom the subject relates, and scenario. We notice a high standard deviation, caused by the fact that some of the scenarios had a low likelihood, possibly because of the random person-meeting combination.

Part III – Projection. In this part, subjects were given pairs of scenarios (from Part II), and they had to select which one they would attend if they could attend only one. We notice that in 69% of the cases, subjects would select the meeting to which they had assigned a higher priority in Part II. This suggests that priority is a good indicator of how people break ties. However, it is not the only thing. We noticed that in most of the cases in which subjects select meetings to which they had assigned a lower priority, those meetings have also a low likelihood. This suggests that when breaking ties between different meetings, subjects also take into account how difficult it would be to reschedule each of the meetings. Also, in this section we see differences between individuals, since there were subjects who consistently chose a certain type of meetings. This can link to the subjects' *personal values* (see also [26,44]). For instance, some subjects consistently picked work meetings or family meetings, which indicates a tie to their value system. This will be explored in future work.

Subjects were also asked about the justification that they would give to the person whose meeting they would cancel. In 89% of the cases, subjects reported that they would give the real reason. Most of the cases in which the subjects would give no reason or a different reason (and not the real one) took place when they chose to attend meetings with a lower priority. Furthermore, many cases involve either not reporting to someone with a higher rank, or not giving details about their meetings with family members.

5 Predicting Priority of Social Situations

In order to address RQ2, we will investigate how to use data from Part II of the user study in order to predict the priority level of social situations based on information about social features. First we will discuss possible options on how to achieve this (Sect. 5.1), and then we will introduce and evaluate our proposed approach (Sect. 5.2).

5.1 Reasoning About Situations

Different strategies can be used to reason about the priority of an event. The most straightforward approach would be to combine the situation cues in an Expected Priority (EP) function, such as:

$$EP = \sum_{f \in \mathcal{F}} w_f v_f$$

where $\mathcal{F}$ is the set of all features considered, and where for all $f \in \mathcal{F}$, v_f refers to the feature value and w_f to the relative weight of feature f in this computation. However, there are two main issues with this approach. First of all, most of the features that we are dealing with have nominal values, so quantifying them is difficult. Furthermore, based on the literature on preference profiles, see e.g., [6], in many decision situations, we hypothesize the weights to be dependent on the individual, making the correct initialization of the weights a challenge.

Another option is to learn a model from our data, and use it to classify new instances. Our proposed approach to do this is to use decision trees [7], because literature suggests that the structure of decision trees is appropriate for reasoning about social relations. First of all, cognitive psychology proposes that social intelligence can have a modular nature [21]. This means different "scripts" are activated in different settings. People recognize these settings from environmental cues, and in turn decide to behave in a certain way. This is similar to the concept of decision trees, in which different combinations of features lead to different decisions. Endsley also suggests that people use different "schemata" to organize and combine knowledge and perceptions in order to comprehend the situation [17]. Moreover, the decision process of decision trees is predictable and transparent. This would allow the agent to *explain* to the user why a certain priority level is assigned to a specific event, which is important since we focus on behaviour support.

Decision trees are graphical representations of a set of rules which can be used to make classifications. Each node of the tree represents a question regarding certain features of the object that is being classified, in this case a social situation, and each branch represents a different answer to that question, in this case the priority level. Nodes below a given node either contain another question, or are given a *label* which assigns a class to the object. The latter are called *leaf-nodes.* Given an object with a set of features and a decision tree, in order to classify the object we traverse the tree until we reach a leaf.

5.2 Model

So far, we have represented the features of the social situations. However, this raw information is not sufficient to draw conclusions about how people evaluate situations. As explained in Sect. 1, in situation awareness literature, this process is called *comprehension* [18]. In this work we explore one general and abstract characteristic of a given situation, namely its *priority*.

As mentioned in the previous section, we will use decision trees to predict priority of social situations. One of the most used methods because of its high accuracy is the Classification and Regression Trees algorithm (*CART*) [7]. CART models are binary trees, which means for every parent node there are two child nodes. Learning a CART model involves selecting features and split points on those features until a suitable tree is constructed. This selection is performed by using a greedy algorithm which minimizes a cost function. We build the model using the R package `rpart` [42]. We use 70% of the data as a training set from which the tree structure was learned, and then test it on the remaining 30%. As a pruning mechanism we limit the maximal depth of the tree to 4.[4]

The learned model is shown in Fig. 3. We remark that, to us, many of the tree splits are intuitive. For instance, the first information that is checked is the setting of the meeting, with casual and sport events on one hand (the left branch) and family and work events on the other (the right branch). This split was to be expected since subjects assigned higher priorities to family and work events.

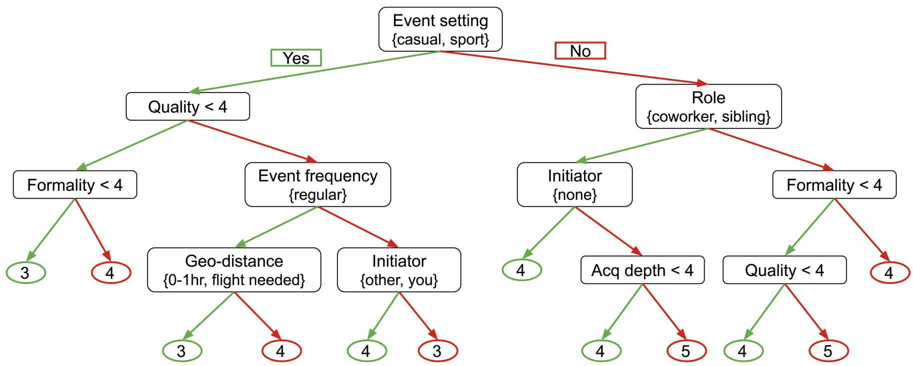

Fig. 3. Decision tree built based on the data. Nodes with categorical features, such as event setting, should be interpreted as "is *event_setting*=`casual OR sport`?"

Since we lack a benchmark in this domain in order to evaluate our model, we compare our result with an algorithm which would predict a random priority (as we offered 5-point scale, chance corresponds to 20%) and with an algorithm which always picks the most selected class, i.e., priority 4, which was selected in 41% of cases. To determine the accuracy of the models, we use the following definition:

[4] The code can be found in: https://github.com/ilir-kola/decisiontree-socialsit.git.

$$accuracy = \frac{Number\_correct\_predictions}{Number\_overall\_predictions}$$

The accuracy of our model on the test set is 47%, thus performing better than the other two algorithms that we used as a benchmark. This means that information about social features can be used to predict priority of a social situation.

6 Discussion and Conclusions

6.1 Research Questions and Hypothesis

Regarding RQ1, in Sect. 4.2 we notice that subjects find our proposed set of features understandable and their quantity appropriate. Furthermore, they find the features relatively expressive. In Sect. 5, we tackle RQ2 by proposing a model which learns a decision tree to predict priority of meetings. We observe that the model performs better than chance, which shows that while this can be a way to predict priority of social situations, more works needs to be done in order to achieve a higher accuracy. This may involve introduction of additional features. The result also contributes in the answer of RQ1, since it suggests that the features allow us to represent social situations in order to learn information about them. Regarding H3, in Sect. 4.2 we see that in 69% of the cases, priority is a good predictor for choosing between overlapping meetings. However, it also shows that it is not the only element, and more dimensions of situations need to be assessed to identify where this difference comes from.

6.2 Contributions

For the benefit of the development of behaviour support agents with social situation awareness, this paper provides the following contributions:

- an upper ontology for representing the salient situation cues and types of features for characterizing dyadic social relationships.
- a set of lower level features which can be used to represent daily life social situations.
- an evaluation of social features via a user study, showing that subjects find the concepts understandable and expressive.
- an evaluation whether decision trees can be used to predict the priority of social situations based on features of social situations, which proved to be the case.

Results presented in this work tend to support the feasibility of our overall approach, but in parallel they open the way for different research questions which need to be explored in more depth in future work.

6.3 Limitations

First of all, the number of people in our user study via which we evaluate our proposed features is relatively small, and the subjects are mostly PhD candidates. This does not allow for a conclusive answer when it comes to understandability and expressiveness among other types of people with, for example, other levels of education. In turn, this also creates limitations when tackling RQ2. First of all, we built the model using a small data set, and learning algorithms need more data in order to generalize better. This is also shown by the high level of overfitting which takes place, as noticed by the fact that the accuracy on the training set is 65%. Moreover, the data is unbalanced, since people mostly give a priority of 4 or 5 to events. The presence of lower priorities would make the evaluation of the algorithm more realistic since we would be able to measure not only the number of correct predictions, but also how far off the incorrect predictions are. The low variance in the data can be explained by the fact that subjects chose people who are very close to them, thus they would prioritize those events.

6.4 Proposed Future Work

Based on the findings reported in this paper, a more extensive experiment can be confidently carried out to obtain a detailed social model that can serve as a background model for behaviour support agents to advise on how to choose between social situations. More data can help not only in building a more accurate model, but also to try out more techniques. Furthermore, that data can also be used to study the correlations between the different features, in order to select a minimal set of features for which to ask the users.

Another interesting approach is to analyze how *personal values* [26,44] affect the way in which subjects think about social situations. Part III of our experiment suggested the existence of individual differences in how people decide which meetings to attend. We will explore whether people with shared personal values make similar choices.

The current model relies fully on information that is acquired directly from the users. In future work, we would like to add sensory data to inform our model. Literature shows that sensory data can be used to perceive social information (e.g., [9]). This line of research would provide useful ways to acquire information without interrupting the user.

Finally, in this work we mostly focus on the modelling of social situations. The next step is to dive deeper into situation comprehension, and reason about different dimensions of social situations (other than priority). Data from the user study suggests that both enjoyment and obligation correlate well with priority, and this correlation is stronger when considering situations in specific settings (enjoyment for casual situations and obligation for work situations). Representing more dimensions of social situations would lead to having a more complete profile of the situation, which in turn enables behaviour support agents to provide more comprehensive help.

Appendix

Part 1

For each person, the following questions were asked:

- What's the name of this person? (e.g. Alice)
- What is the role of Alice towards you? options: {`partner, parent, sibling, child, friend, extended family member, neighbor, coworker, supervisor, member of the same group (e.g., sports team), other`}
- What's the hierarchy rank (from a formal point of view) of Alice towards you? options: {`higher, lower, same, n.a.`}
- How would you consider the quality of your relationship with Alice? options: $\texttt{Likert}_5$
- What's the geographical distance between you and Alice? options: {`0-1 h, 1-2 h, 2-4 h, I would need to take a flight`}
- How well do you know Alice? options: $\texttt{Likert}_5$
- How often are you in touch with Alice? options: $\texttt{Likert}_5$
- How much do you trust Alice? options: $\texttt{Likert}_5$
- How formal is your relationship with Alice? options: $\texttt{Likert}_5$

Part 2

For each scenario, the following questions were asked. For all, the answer option was a 5-point Likert scale:

- What priority would you assign to this meeting?
- What priority do you think the other person would assign to this meeting?
- To what extent would you feel obligated to attend this meeting?
- To what extent do you think the other person would feel obligated to attend this meeting?
- To what extent would you enjoy attending this meeting?
- To what extent do you think the other person would enjoy attending this meeting?
- How likely are you to encounter this scenario in your life?

Part 3

Two scenarios were chosen randomly and shown to the subject, and the following questions were asked:

- If they were planned to happen at the same time, which of the two scenarios would you attend? options: {`Scenario 1, Scenario 2`}
- What explanation would you give to the person whose meeting you would have to cancel? options: {`no explanation, the real reason, some other reason`}.

References

1. Ajmeri, N., Guo, H., Murukannaiah, P.K., Singh, M.P.: Designing ethical personal agents. IEEE Internet Comput. **22**(2), 16–22 (2018)
2. Ajmeri, N., Murukannaiah, P.K., Guo, H., Singh, M.P.: Arnor: modeling social intelligence via norms to engineer privacy-aware personal agents. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, pp. 230–238. International Foundation for Autonomous Agents and Multiagent Systems (2017)
3. Antonucci, T.C., Ajrouch, K.J., Birditt, K.S.: The convoy model: explaining social relations from a multidisciplinary perspective. Gerontologist **54**(1), 82–92 (2013)
4. Antonucci, T.C., Akiyama, H.: Social networks in adult life and a preliminary examination of the convoy model. J. Gerontol. **42**(5), 519–527 (1987)
5. Barwise, J., Perry, J.: Situations and attitudes. J. Philos. **78**(11), 668–691 (1981)
6. Boutilier, C.: A POMDP formulation of preference elicitation problems. In: Eighteenth National Conference on Artificial Intelligence, pp. 239–246. American Association for Artificial Intelligence, Menlo Park (2002). http://dl.acm.org/citation.cfm?id=777092.777132
7. Breiman, L., Friedman, J., Olshen, R., Stone, C.: Classification and Regression Trees. Routledge, Abingdon (1994)
8. Burt, R.S.: Network items and the general social survey. Soc. Netw. **6**(4), 293–339 (1984)
9. Cabrera-Quiros, L., Demetriou, A., Gedik, E., van der Meij, L., Hung, H.: The MatchNMingle dataset: a novel multi-sensor resource for the analysis of social interactions and group dynamics in-the-wild during free-standing conversations and speed dates. IEEE Trans. Affect. Comput. (2018). https://doi.org/10.1109/TAFFC.2018.2848914
10. Devlin, K.: Logic and Information. Cambridge University Press, Cambridge (1995)
11. Dignum, F.: Interactions as social practices: towards a formalization. arXiv preprint arXiv:1809.08751 (2018)
12. Dignum, F., Prada, R., Hofstede, G.J.: From autistic to social agents. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems, pp. 1161–1164. International Foundation for Autonomous Agents and Multiagent Systems (2014)
13. Dignum, V.: A Model for Organizational Interaction: Based on Agents, Founded in Logic. SIKS (2004)
14. Dignum, V., Dignum, F.: Contextualized planning using social practices. In: Ghose, A., Oren, N., Telang, P., Thangarajah, J. (eds.) COIN 2014. LNCS (LNAI), vol. 9372, pp. 36–52. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25420-3_3
15. Dudzik, B., et al.: Context in human emotion perception for automatic affect detection: a survey of audiovisual databases. In: Proceedings of the 8th International Conference on Affective Computing & Intelligent Interaction. Association for the Advancement of Affective Computing (2019)
16. d'Inverno, M., Luck, M., Noriega, P., Rodriguez-Aguilar, J.A., Sierra, C.: Weaving a fabric of socially aware agents. In: Kinny, D., Hsu, J.Y., Governatori, G., Ghose, A.K. (eds.) PRIMA 2011. LNCS (LNAI), vol. 7047, pp. 263–274. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25044-6_21
17. Endsley, M.R.: Toward a theory of situation awareness in dynamic systems. Hum. Factors **37**(1), 32–64 (1995)

18. Endsley, M.R.: Theoretical underpinnings of situation awareness: a critical review. Situat. Aware. Anal. Meas. **1**, 24 (2000)
19. Fiske, A.P.: The four elementary forms of sociality: framework for a unified theory of social relations. Psychol. Rev. **99**(4), 689 (1992)
20. Fullam, K.K., Klos, T., Muller, G., Sabater-Mir, J., Barber, K.S., Vercouter, L.: The agent reputation and trust (ART) testbed. In: Stølen, K., Winsborough, W.H., Martinelli, F., Massacci, F. (eds.) iTrust 2006. LNCS, vol. 3986, pp. 439–442. Springer, Heidelberg (2006). https://doi.org/10.1007/11755593_32
21. Gigerenzer, G.: The modularity of social intelligence. In: Machiavellian Intelligence II: Extensions and Evaluations, vol. 2, no. 264, pp. 264–288 (1997)
22. Granovetter, M.: The strength of weak ties: a network theory revisited. In: Sociological Theory pp. 201–233 (1983)
23. Granovetter, M.: Weak ties and strong ties. Am. J. Sociol. **78**, 1360–1380 (1973)
24. Hoekstra, R.: Ontology representation design patterns and ontologies that make sense. In: Proceedings of the 2009 Conference on Ontology Representation: Design Patterns and Ontologies that Make Sense, pp. 1–236. IOS Press (2009)
25. Kahn, R., Antonucci, T.: Convoys over the life course: attachment, roles, and social support. In: Life-Span Development and Behavior (1980)
26. Kayal, A., Brinkman, W.P., Neerincx, M.A., Riemsdijk, M.B.V.: Automatic resolution of normative conflicts in supportive technology based on user values. ACM Trans. Internet Technol. **18**(4), 41:1–41:21 (2018)
27. Kokar, M.M., Matheus, C.J., Baclawski, K.: Ontology-based situation awareness. Inf. Fusion **10**(1), 83–98 (2009)
28. Kola, I., Jonker, C.M., van Riemsdijk, M.B.: Modelling the social environment: towards socially adaptive electronic partners. In: 10th International Workshop on Modelling and Reasoning in Context (2018)
29. Kola, I., Jonker, C.M., van Riemsdijk, M.B.: Pilot experiment exploring the priority of social situations (2019). https://doi.org/10.4121/uuid:e18fb318-c1d4-4ccc-9b4f-be48e1ee49e2
30. Mainela, T.: Types and functions of social relationships in the organizing of an international joint venture. Ind. Mark. Manag. **36**(1), 87–98 (2007)
31. Murukannaiah, P.K., Singh, M.P.: Platys social: relating shared places and private social circles. IEEE Internet Comput. **16**(3), 53–59 (2012)
32. Myers, K.L., Yorke-Smith, N.: Proactivity in an intentionally helpful personal assistive agent. In: AAAI Spring Symposium: Intentions in Intelligent Systems, pp. 34–37 (2007)
33. Parrigon, S., Woo, S.E., Tay, L., Wang, T.: Caption-ing the situation: a lexically-derived taxonomy of psychological situation characteristics. J. Pers. Soc. Psychol. **112**(4), 642 (2017)
34. Pereira, G., Prada, R., Santos, P.A.: Integrating social power into the decision-making of cognitive agents. Artif. Intell. **241**, 1–44 (2016)
35. Pervin, L.A.: A free-response description approach to the analysis of person-situation interaction. ETS Res. Bull. Ser. **1975**(2), i-26 (1975)
36. Pinder, C., Vermeulen, J., Cowan, B.R., Beale, R.: Digital behaviour change interventions to break and form habits. ACM Trans. Comput.-Hum. Interact. (TOCHI) **25**(3), 15 (2018)
37. Rauthmann, J.F., et al.: The situational eight diamonds: a taxonomy of major dimensions of situation characteristics. J. Pers. Soc. Psychol. **107**(4), 677 (2014)
38. Ring, P.S., Van de Ven, A.H.: Developmental processes of cooperative interorganizational relationships. Acad. Manag. Rev. **19**(1), 90–118 (1994)

39. Saucier, G., Bel-Bahar, T., Fernandez, C.: What modifies the expression of personality tendencies? Defining basic domains of situation variables. J. Pers. **75**(3), 479–504 (2007)
40. Singh, M.P.: Norms as a basis for governing sociotechnical systems. ACM Trans. Intell. Syst. Technol. (TIST) **5**(1), 21 (2013)
41. Tambe, M., Bowring, E., Pearce, J.P., Varakantham, P., Scerri, P., Pynadeth, D.V.: Electric elves: what went wrong and why. AI Mag. **29**, 23 (2008)
42. Therneau, T., Atkinson, B.: rpart: Recursive Partitioning and Regression Trees (2018). https://CRAN.R-project.org/package=rpart. R Package Version 4.1-13
43. Tielman, M., Brinkman, W.-P., Neerincx, M.A.: Design guidelines for a virtual coach for post-traumatic stress disorder patients. In: Bickmore, T., Marsella, S., Sidner, C. (eds.) IVA 2014. LNCS (LNAI), vol. 8637, pp. 434–437. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09767-1_54
44. Tielman, M.L., Jonker, C.M., van Riemsdijk, M.B.: What should i do? Deriving norms from actions, values and context. In: Proceedings of the 10th International Workshop on Modelling and Reasoning in Context, MRC 2018, pp. 35–40. CEUR Workshop Proceedings (2018)
45. Uzzi, B.: Social structure and competition in interfirm networks: the paradox of embeddedness. Adm. Sci. Q. **42**, 35–67 (1997)
46. Van Riemsdijk, M.B., Jonker, C.M., Lesser, V.: Creating socially adaptive electronic partners: interaction, reasoning and ethical challenges. In: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, pp. 1201–1206. International Foundation for Autonomous Agents and Multiagent Systems (2015)
47. Williamson, O.E.: Markets and Hierarchies: Analysis and Antitrust Implications. Free Press, New York (1975)
48. Yau, S.S., Liu, J.: Hierarchical situation modeling and reasoning for pervasive computing. In: The Fourth IEEE Workshop on Software Technologies for Future Embedded and Ubiquitous Systems, p. 6. IEEE (2006)
49. Yolum, P., Singh, M.P.: Service graphs for building trust. In: Meersman, R., Tari, Z. (eds.) OTM 2004. LNCS, vol. 3290, pp. 509–525. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30468-5_32
50. Zavala, L., et al.: Platys: from position to place-oriented mobile computing. AI Mag. **36**(2), 50–62 (2015)
51. Zhao, R., Papangelis, A., Cassell, J.: Towards a dyadic computational model of rapport management for human-virtual agent interaction. In: Bickmore, T., Marsella, S., Sidner, C. (eds.) IVA 2014. LNCS (LNAI), vol. 8637, pp. 514–527. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09767-1_62

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.