FEATURE ARTICLE: COGNITIVE AGENTS

Toward Social Situation Awareness in Support Agents

Ilir Kola and Pradeep K. Murukannaiah, *Delft University of Technology, Delft, 2628 CD, The Netherlands,*
Catholijn M. Jonker, *Delft University of Technology, Delft, 2628 CD, The Netherlands, and also University of Leiden, Leiden, 2311 EZ, The Netherlands*
M. Birna van Riemsdijk, *University of Twente, Enschede, 7522 NB, The Netherlands*

Artificial agents that support people in their daily activities (e.g., virtual coaches and personal assistants) are increasingly prevalent. Since many daily activities are social in nature, support agents should understand a user's social situation to offer comprehensive support. However, there are no systematic approaches for developing support agents that are social situation aware. We identify key requirements for a support agent to be social situation aware and propose steps to realize those requirements. These steps are presented through a conceptual architecture centered on two key ideas: 1) conceptualizing social situation awareness as an instantiation of "general" situation awareness, and 2) using situation taxonomies for such instantiation. This enables support agents to represent a user's social situation, comprehend its meaning, and assess its impact on the user's behavior. We discuss empirical results supporting the effectiveness of the proposed approach and illustrate how the architecture can be used in support agents through two use cases.

Human behavior is a function of a person's characteristics as well as the situation.[1] Thus, to support people in their daily lives, artificial agents must represent and reason about not only the personal characteristics but also the situation of a user. To take a user's situation into account, support agents should reason about the user's surrounding entities and how they relate to each other and the user.

Most of our daily situations are social in nature. We collaborate with co-workers, spend weekends with family and friends, and share most of our moments with people. Thus, support agents should account for this social dimension of our lives.

We define social situation awareness and propose the building blocks necessary for support agents to be social situation aware. These building blocks are presented through a conceptual architecture inspired by work on "general" situation awareness,[2] which we instantiate with concepts from social sciences[3] to account for the requirements of modeling social situations. This work serves as a proof of concept showing that the building blocks of social situation awareness can be implemented in support agents and discusses the remaining steps for successful deployment of a full-fledged agent.

WHAT IS SOCIAL SITUATION AWARENESS?

Yang *et al.*[4] defined a situation as "a combination of the individually interpreted, implicit, and unique understandings, and the culturally shared, explicit, and common understandings of the surroundings that produce and constrain human behavior." We define a *social situation* as a type of situation that involves more than one person. Thus, a social situation involves not only the typical situational elements, such as time and place but also social elements, such as the quality of the relationships and contact frequency between the user and other people in the situation.

The social elements of a situation influence user behavior. For instance, consider two situations: one in which a user has dinner with a friend and another in which the user has dinner with a prospective employer.

1541-1672 © 2022 IEEE
Digital Object Identifier 10.1109/MIS.2022.3163625

SITUATION TAXONOMIES

Situations are abstract entities, which makes assigning meaning to them challenging. Studies in social psychology[3] suggest that people interpret situations by creating impressions of them as if they were real entities that have specific *psychological characteristics*. Understanding situations in terms of these characteristics allows people to better navigate the world by using these characteristics to predict what will happen and coordinate behavior accordingly. We propose that support agents should similarly treat situations as real entities with psychological characteristics.

Psychological characteristics provide a high-level subjective interpretation of situations and are widely studied, and different taxonomies have been developed. Here, we present the elements of the DIAMONDS taxonomy.[3] We choose this taxonomy because it is designed to cover daily situations and it offers a validated scale for measuring the psychological characteristics of situations. The taxonomy comprises the following characteristics in terms of which situations can be described:

- *Duty*—situations where a job has to be done, minor details are important, and rational thinking is called for;
- *Intellect*—situations that afford an opportunity to demonstrate intellectual capacity;
- *Adversity*—situations where you or someone else are (potentially) being criticized, blamed, or under threat;
- *Mating*—situations where potential romantic partners are present, and physical attractiveness is relevant;
- *Positivity*—playful and enjoyable situations, which are simple and clear cut;
- *Negativity*—stressful, frustrating, and anxiety-inducing situations;
- *Deception*—situations where someone might be deceitful. These situations may cause feelings of hostility;
- *Sociality*—situations where social interaction is possible, and close personal relationships are present or have the potential to develop.

Rauthmann *et al.*[3] suggested that people use these psychological characteristics to ascribe meaning to situations. Furthermore, they show that psychological characteristics of situations correlate with various situation cues, as well as behaviors that people exhibit in those situations. For instance, a high level of duty is characteristic of work situations, and typical behaviors for situations with a high level of duty are concentrating and displaying ambition. This corresponds to our definition of social situation awareness: psychological characteristics of situations can be used for social situation comprehension, and are related to both social situation perception and social situation projection.

In these two situations, despite similar environmental elements, the user's behavior can be different because of the different relationships among the people in these situations.

Endsley[2] described a prominent model of situation awareness consisting of three levels:

1) *perception:* representing the status, attributes, and dynamics of relevant elements in the environment;
2) *comprehension:* representing a higher level synthesized meaning of the elements of the environment; and
3) *projection:* representing the ability to project the future status of the elements of the environment.

Adapting Endsley's model, we define *social situation awareness* as:

> a support agent's ability to perceive the social elements of a situation, to comprehend their meaning, and to infer their effect on the behavior of the user.

REQUIREMENTS FOR SOCIAL SITUATION-AWARE AGENTS

Different context awareness architectures have been proposed for different purposes. Alegre *et al.*[5] provided a review of existing approaches, and suggest that one

of the reasons for the variety of existing approaches is the need for specific architectures in each domain. However, none of the reviewed approaches tackles social situations specifically. Our research fills this gap. Focusing on social situations motivates us to take into account the human aspects of a situation as opposed to the technical aspects investigated in related work, such as geo-spatial locations and other physical elements of context. Furthermore, the focus of existing approaches is on information that can be acquired through sensors, which is processed to determine actions that are occurring in the environment. Our work complements these approaches by focusing on the psychological meaning of situations. Based on these differences, we formulate the following requirements for support agents to be social situation aware.

Combining sensory data with a user's mental constructs: Perceiving social situations relies not only on information that can be detected through sensors, but also on a user's mental constructs. For instance, in a situation where a user is meeting another person for dinner, it is difficult to detect the features of their relationship from sensors alone. This information can be important, e.g., a dinner with a friend is very different from a dinner with a potential employer. Therefore, the agent needs to be able to elicit information about the user's mental constructs, such as social relations, which may not be available via sensors.

Variety of social situations: A flexible support agent should be able to represent a wide variety of social situations a user may encounter. To do so, an agent must identify a variety of social dimensions characterizing a situation. Furthermore, the agent should be able to interpret this situation variety by translating social features into abstractions to determine appropriate support, e.g., using prespecified rules to categorize situations into a limited number of higher level situations, or using machine learning to derive information that can be used for reasoning about support.

Interpreting the meaning of situations: Existing work on situation awareness addresses the comprehension step by determining how the perceived objects in a situation are interrelated[6] and recognizing the situation type. For instance, if two users are perceived to be in the same office, the comprehension step would say that the user is in a meeting. However, in social situations, knowing the type of situation is not sufficient to determine the type of support needed since it is possible to infer different meanings from this information. For example, being in a meeting with a supervisor is different from a meeting with a potential client. Support agents need to be able to distinguish the different meanings of such social situations.

Value-aware support: Agents should provide support that is consistent with the user's goals and preferences. In social science, it has been argued that the essence of a situation is its affordance of human goals and motives.[3] A way to represent human motives are personal values. Values, such as independence or success, which express what people find important in life have been found to be key drivers of human decisions, and value preferences exhibit cross-situational consistency.[7] Since providing support in social situations is ultimately about aligning with the user's underlying motivations, we suggest the use of values for personalization.

Explainability and directability: Support agents need to be able to explain their suggestions to users. For instance, consider an agent that supports a healthy lifestyle. If the agent merely suggests to the user to avoid going to a party, this advice might be unexpected. However, if the agent informs the user that going to parties usually leads to smoking, which demotes the value of "health," then the user can make an informed decision. Furthermore, the user should be able to direct the agent on how to act. Continuing our example, the user should be able to inform the agent that since the party is in a non-smoking venue, health would not be demoted. The agent can then use this information in future situations.

Although variants of these requirements are mentioned in existing work, to the best of our knowledge our formulation and approach toward tackling them in an integrated way is new. The key novel elements in our requirements are the consideration of how to ascribe meaning to social situations, the emphasis on user interaction, and a hybrid human–machine approach for social situation awareness and support.

CONCEPTUAL ARCHITECTURE FOR SOCIAL SITUATION AWARENESS

We identify the core elements, and their interrelations, for creating social situation aware agents by presenting a conceptual architecture (Figure 1). The architecture consists of two main components: a social situation awareness component, and a user interaction component. The first is an instantiation of the three-level situation awareness model proposed by Endsley[2] with social concepts. The second comprises interaction modules needed for integrating situation awareness reasoning with the supportive function of the agent. We provide directions for implementing these components.

Level 1: Social Situation Perception

The goal of the perception level is to obtain a representation of the salient aspects of a social situation. This

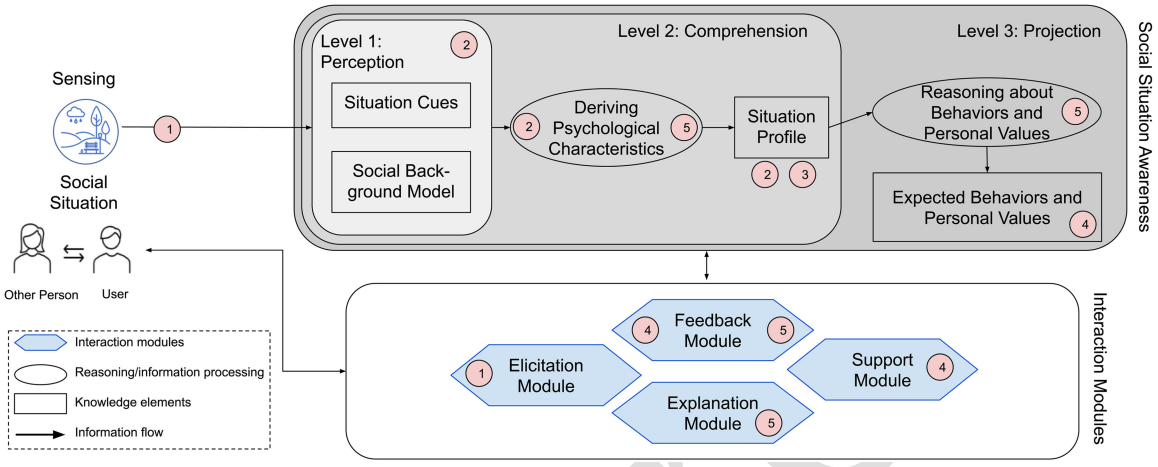

FIGURE 1. Architecture of a social situation aware support agent. The numbers in red circles represent the requirements (see Table 1) that the elements of the architecture address.

information can come from sensory data and interaction with the user. To account for a wide variety of situations, the information included in this level should allow representing arbitrary social situations. Kola *et al.*[8] proposed an approach to model arbitrary social situations through a two-level ontology distinguishing *situation cues* and social relationship features (*social background model*). Rosatelli *et al.*[9] proposed an approach where data from wearable sensors is processed with deep learning techniques to assess information, such as roles in social interactions.

Level 2: Social Situation Comprehension

In this level, the perceived elements are used to infer a social situation profile, characterizing the situation along meaningful dimensions.

Knowledge Elements

To describe the meaning of a social situation, we propose to use the psychological characteristics of situations (see section on situation taxonomies). The idea is to describe each social situation by a set of features (the *situation profile*) that represent the psychological characteristics of that situation. These characteristics describe a user's subjective understanding of a situation. A key advantage of this approach is that it offers a fixed number of dimensions based on which it is possible to represent and compare different situations.

Reasoning

To determine the psychological characteristics of a situation, one may follow a rule-based or a machine learning approach. A rule-based approach provides explicit reasoning, but requires extensive design time specifications. A machine learning approach supports situation variety, e.g., by offering predictions for unseen examples, but offers limited explainability.

Level 3: Social Situation Projection

In this level, the agent uses the situation profile to predict how the user is likely to behave in a social situation, and what values are affected.

Knowledge Elements

In the classic situation awareness model, the projection level captures how the situation develops as a whole. To fulfill the personalization requirement, we propose that in the projection level the agent needs to predict what behavior the user is likely to exhibit, and the personal values that are promoted or demoted in a given situation. The former allows the agent to provide proactive support, and the latter enables the agent to help the users in a value-aligned manner.

Reasoning

This component takes the situation profiles as input, and predicts the expected behavior and the promoted and demoted values. A possible way to do this is by grouping similar situations based on their profile, and studying the patterns of behavior and values in each group of situations, as done by Kola *et al.*[10]

Interaction Modules

An agent needs to interact with the user to give and acquire information necessary for support. We foresee the need for four interaction modules. In this article,

we focus on describing the role of these modules as part of the envisaged support agent. In order to realize the interaction modules and create a full-fledged social situation-aware agent, open research challenges regarding human–machine hybrid intelligence[11] need to be addressed.

Elicitation Module

The elicitation module interacts with the user to elicit necessary information that cannot be acquired from a sensor. User interaction is needed during both initialization and run time. During initialization, the goal is to gather information that remains relatively stable, e.g., information about a user's social relationships with their most frequent contacts, needed to form the social background model. This ensures that for most social situations that the user encounters, the social background model already contains the needed information, thus avoiding overloading the user with information requests after initialization. During run time, the module detects when certain information is missing regarding a specific social situation, e.g., the role of the other person, and asks the user. The Platys framework[12] can be used to reduce the possible burden of information elicitation for the user. Platys employs an active learning approach, which asks a user to provide context information only if the predictions with existing sensor readings are uncertain, which reduces the overall effort for the user.

Support Module

After going through the social situation awareness levels, an agent can reason about the support it can provide. One of the proposed requirements is for the agent to personalize support according to the needs and the values of the user. This information can be contained in a user model within the support module. The support module can then compare the user preferences with the information coming from Level 3 of the architecture regarding expected user behavior and values. Support is needed when there is a mismatch between the preferred and the expected behavior of the user in a situation, or when the situation affects a value important to the user.

Explanation Module

To make an agent's support actions explainable, we propose to use meaningful social notions in each level of the architecture, derived through explainable reasoning and learning techniques. An advantage of a multilevel architecture is that explanations can be given on different levels: the agent can: 1) give insight on the suggestion relating it to a certain personal value or preferred behavior (Level 3); 2) explain why a certain behavior or personal value is expected in a specific social situation by referring to the psychological characteristics of that situation (Level 2); and 3) give further insight on the situation cues and social relationship aspects that cause the situation to have those specific psychological characteristics (Level 1).

Feedback Module

It should be possible for the user to notify the agent when a support action or its explanation is not satisfactory. The feedback module achieves this by interacting with the user to determine whether there has been a mistake in one of the reasoning steps or whether some information in the knowledge base needs to be updated. The agent can then integrate this feedback into its reasoning mechanisms and knowledge bases at the appropriate level. How exactly such updates are to be performed and represented is an open research question.

EMPIRICAL EVIDENCE

In this section, we present empirical evidence that supports our proposed three-level social situation awareness component. The social situation awareness component is an instantiation of the well-known model of situation awareness by Endsley.[2] The model's diverse applications suggest that the three-level approach as a whole is beneficial.

In past work, the different levels of the social situation awareness were implemented and evaluated through human-grounded studies.[14] Human-grounded evaluations involve real people who are presented with simplified tasks, and are particularly useful in cases, such as ours, where the goal is to evaluate reasoning components and a full-fledged agent cannot yet be implemented due to open challenges in interaction modules. In Kola *et al.*,[13] we showed that transitioning through the three levels of the architecture is possible: using data collected from a large user study, we presented an approach in which it is possible to predict Level 2 information from Level 1 inputs, and then in turn use the predicted Level 2 information as input for predicting Level 3 information. Furthermore, we showed how Levels 1 and 2 information can be used as a basis for creating explanations that are satisfying for people. In this section, we give details on how the different levels of the social situation awareness module have been successfully implemented and evaluated in generic domains, e.g., to assess the promoted or demoted personal values of a social situation, or specific domains, such as reasoning about

TABLE 1. Key requirements and how they are addressed in our proposed approach.

Requirement	How it is addressed	Empirical evidence
1) Combining sensory data with mental constructs of the user	Perception level based on sensory data and user-elicited information	[8]
2) Variety of situations	Use concepts from the social sciences to allow representing arbitrary situations; Use machine learning to learn connections between Levels 1 and 2	[3,8]
3) Interpreting the meaning of situations	Derive the psychological characteristics of situations	[3,13]
4) Value-aware support	Base support on the personal values of the user; Have feedback module which allows personalization	[10]
5) Explainability and directability	Use explainable techniques; Explanation module techniques Feedback module	[13]

the priority of social situations. Furthermore, in Table 1 we present how each architectural element tackles the identified requirements.

Level 1

In Kola *et al.*,[8] we proposed an ontology to tackle the perception level. The ontology models situation cues, describing the situation, and social relationship features, describing the relationship of the user with the people in the situation. We evaluated this approach via a user study in which participants were asked about their social relationships using the features proposed in the ontology. Participants considered the ontology to contain an appropriate amount of information (average answer=3, SD=0.61 on a five-points Likert scale where 1=too little information, 3=appropriate information, 5=too much information) and to be fairly representative of their social relationships (average answer=3, SD=0.79 on a five-points Likert scale where 1=very little representative and 5=very much representative). This study suggests that it is possible to have a model of a social situation that includes a user's mental constructs, in particular describing social aspects of situations, thus fulfilling Requirement 1.

Level 2

In this level, we suggest ascribing meaning to situations through the psychological characteristics. Rauthmann *et al.*[3] conducted validation studies involving hundreds of participants across different countries and cultures, showing that the DIAMONDS taxonomy can be used to give meaning to arbitrary situations, thus providing evidence for Requirements 2 and 3. A technical requirement of the architecture is the ability to derive these psychological characteristics from the information from Level 1. To investigate this, we collected Levels 1 and 2 data through a crowdsourcing user study.[13] Using this data, we showed that machine learning models can be created that predict psychological characteristics of situations from Level 1 information with an average error of 1.14 on a six-point Likert scale, outperforming benchmark results.

Level 3

In Kola *et al.*,[10] we proposed an approach that groups situations based on psychological characteristics and shows that different personal values are promoted or demoted in specific groups of situations. For instance, we noticed that situations with high intellect and duty promote the values, of helpfulness and capability. This helps fulfill Requirement 4. Furthermore, this shows that transition from Levels 2 to 3 is feasible with respect to personal values. To show that this transition is also possible in terms of expected behaviors, in Kola *et al.*'s work,[13] we used psychological characteristics of situations as input to predict expected user behavior regarding social priorities with an error of 0.98 on a seven-point Likert scale for actual values of the characteristics, and with an error of 1.37 for predicted values of psychological characteristics based on Level 1 information.

USE CASES

We illustrate how the components of our approach could be included in intelligent agents that provide support via two use cases: 1) agenda management[16] and 2) location data sharing[15] support agents. Although these use cases are quite different, the high-level components of our approach can be instantiated for each use case, as shown in Table 2. This illustrates how our approach can serve as a blueprint for including social situation awareness in support agents.

Agenda Management Support Agent

Kola *et al.*[16] introduced an agenda management support agent, whose goal is to assess a user's priorities and make suggestions based on the priority levels when different meetings overlap. Table 2 illustrates the information modeled in the different components of the

TABLE 2. Concepts that can be modelled and role of modules in two use cases.

Use case	Agenda management support agent	Value-based location sharing agent
Level 1	Social background features of other person (e.g., role, hierarchy level)	Location-related features; Other people present
Level 2	Psychological characteristics of situation (e.g., duty, intellect)	Psychological characteristics of situation (e.g., sociality)
Level 3	Predict priority of meetings	Assess how values are affected
Elicitation	Social situation features	Personal values
Support	Suggest which meeting to attend based on priority	Provide value-aligned support
Explanation	Why a meeting was suggested	Why a location was shared with someone
Feedback	Adapt priority prediction model	Adapt value assessment

architecture. Level 1 (*perception*) includes information such as the role of the other person and their hierarchy level. The agent uses the perceived information to assess the psychological characteristics of the situation, which are modeled through the DIAMONDS taxonomy (*comprehension*). From this information, the agent determines a priority level for every social situation (*projection*). If two meetings overlap, the agent suggests to the user to attend the meeting with higher priority and reschedule the other. The user can ask for the reason behind the suggestion, and explanations can be given based on Levels 1 or 2 information. If the user does not accept the suggestion, the *feedback* module asks about the reasons and incorporates the feedback into the knowledge base and reasoning processes to better predict priority in similar future situations.

In this use case, following our approach allows us to explicitly take into account social aspects of the situation, which are modeled from the point of the view of the user through the elicitation module and the perception level. Furthermore, our proposed situation comprehension approach allows for a richer representation and understanding of situations, which in turn allows us to better assess their priority. For instance, an agent may determine that meetings involving a high level of duty are more important for a specific user.

Value-Based Location-Sharing Agent

Kayal *et al.*[15] proposed a model for choosing among conflicting agreements about social sharing of location data based on the users' personal values. They show that an agent can help in resolving conflicting commitments by knowing the value preferences of the user and the promoted values of different location-sharing commitments.

Our social situation awareness framework can extend this approach. Level 3 (*projection*) enables the agent to automatically assess which values are promoted or demoted in a situation. Once this information is available, the *support* module can assess whether a specific location sharing activity is aligned with the values of the user. Including information about social relations (*perception*) allows a prediction of values based on a richer model. Furthermore, explicitly modeling the psychological characteristics of the situation (*comprehension*) can be beneficial since these have been shown to be a good predictor of personal values afforded in a situation.[3,10] For instance, the agent may infer that situations taking place in specific locations involve a high level of sociality, and such situations also tend to promote the value social recognition. If the value is important to the user, the agent would share the location data. This information would also facilitate *explanations*: if the user asks why the location was shared, the agent would explain that it had inferred that the situation promotes social recognition because it involves a high level of sociality. If this inference is not correct, the *feedback* module will adapt the value assessment model accordingly.

CONCLUSION

We outline the elements needed for social situation awareness in support agents and illustrate their practical benefits. Existing work (e.g., Kola *et al.*'s work[10,13,16]) has shown promising results in implementing the different levels of social situation awareness, as well as in automatically transitioning between the levels using data from studies conducted with real people. This work serves as a proof of concept for social situation awareness in support agents. However, more research from

different communities is needed to go from this proof of concept to a full-fledged agent. First, the interactive modules will have a crucial role in the successful implementation of an agent that can be tested on real tasks with users. Realizing these requires further research investigating how hybrid intelligent systems can be made collaborative, adaptive, responsible, and explainable.[11] This includes advances in integrating active learning approaches in order to better personalize the prediction models for specific users based on their feedback. Finally, this proposed architecture should be integrated with work on interpreting the meaning of social signals such as body language in social interactions.[9] This would allow the agent to take into account the dynamics of a social situation as it unfolds, allowing it to integrate social situation understanding based on social relations with observed social signals.

ACKNOWLEDGMENTS

This work is part of the research programme the Core-SAEP, with project number 639.022.416, which is financed by the Netherlands Organisation for Scientific Research.

REFERENCES

1. K. Lewin, "Field theory and experiment in social psychology: Concepts and methods," *Amer. J. Sociol.*, vol. 44, no. 6, pp. 868–896, 1939.
2. M. R. Endsley, "Toward a theory of situation awareness in dynamic systems," *Hum. Factors*, vol. 37, no. 1, pp. 32–64, 1995.
3. J. F. Rauthmann *et al.*, "The situational eight diamonds: A taxonomy of major dimensions of situation characteristics," *J. Pers. Social Psychol.*, vol. 107, no. 4, pp. 677–718, 2014.
4. Y. Yang, S. J. Read, and L. C. Miller, "The concept of situations," *Social Pers. Psychol. Compass*, vol. 3, no. 6, pp. 1018–1037, 2009.
5. U. Alegre, J. C. Augusto, and T. Clark, "Engineering context-aware systems and applications: A survey," *J. Syst. Softw.*, vol. 117, pp. 55–83, 2016.
6. N. Baumgartner and W. Retschitzegger, "A survey of upper ontologies for situation awareness," in *Proc. 4th IASTED Int. Conf. Knowl. Sharing Collaborative Eng.*, 2006, pp. 1–9.
7. S. H. Schwartz, "Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries," *Adv. Exp. Social Psychol.*, vol. 25, no. 1, pp. 1–65, 1992.
8. I. Kola, C. M. Jonker, and M. B. van Riemsdijk, "Who's that? social situation awareness for behaviour support agents," in *Proc. Int. Workshop Eng. Multi-Agent Syst.*, 2019, pp. 127–151.
9. A. Rosatelli, E. Gedik, and H. Hung, "Detecting f-formations & roles in crowded social scenes with wearables: Combining proxemics & dynamics using lstms," in *Proc. 8th Int. Conf. Affect. Comput. Intell. Interact. Workshops Demos*, 2019, pp. 147–153.
10. I. Kola, C. M. Jonker, M. L. Tielman, and M. B. van Riemsdijk, "Grouping situations based on their psychological characteristics gives insight into personal values," in *Proc. 11th Int. Workshop Model. Reasoning Context*, 2020, pp. 17–26.
11. Z. Akata *et al.*, "A research agenda for hybrid intelligence: Augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence," *Computer*, vol. 53, no. 8, pp. 18–28, 2020.
12. P. K. Murukannaiah and M. P. Singh, "Platys: An active learning framework for place-aware application development and its evaluation," *ACM Trans. Softw. Eng. Methodol.*, vol. 24, no. 3, pp. 1–32, 2015.
13. I. Kola, C. M. Jonker, and M. B. van Riemsdijk, "Using psychological characteristics of situations for social situation comprehension in support agents," 2021, *arXiv:2110.09397*.
14. F. Doshi-Velez and B. Kim, "Considerations for evaluation and generalization in interpretable machine learning," in *Explainable and Interpretable Models in Computer Vision and Machine Learning*. H. J. Escalante, S. Escalera, I. Guyon, X. Baró, Y. Güçlütürk, U. Güçlü, and M. van Gerven, Eds., New York, NY, USA: Springer, 2018, pp. 3–17.
15. A. Kayal, W.-P. Brinkman, M. A. Neerincx, and M. B. V. Riemsdijk, "Automatic resolution of normative conflicts in supportive technology based on user values," *ACM Trans. Internet Technol.*, vol. 18, no. 4, pp. 1–21, 2018.
16. I. Kola, M. L. Tielman, C. M. Jonker, and M. B. van Riemsdijk, "Predicting the priority of social situations for personal assistant agents," in *Proc. Int. Conf. Princ. Pract. Multi-Agent Syst.*, 2020, pp. 231–247.

ILIR KOLA is a Ph.D. candidate in AI at TU Delft, Delft, 2628 CD, The Netherlands. His research interests include how to enable supportive agents take into account the effect of the social environment on the behavior of the user. Kola received his M.Sc. degree in computer science from the University of Freiburg, Freiburg im Breisgau, Germany. Contact him at i.kola@tudelft.nl.

PRADEEP K. MURUKANNAIAH is an assistant professor (with Tenure) in interactive intelligence at TU Delft, Delft, 2628 CD, The Netherlands. His research focuses on engineering of hybrid intelligence systems that involve collaborations among humans and artificial agents. Murukannaiah received his Ph.D. degree in computer science from North Carolina State University, Raleigh, NC, USA. Contact him at p.k.murukannaiah@tudelft.nl.

CATHOLIJN M. JONKER is a full professor in interactive intelligence at TU Delft, Delft, 2628 CD, The Netherlands, and a full professor of explainable AI at Leiden University, Leiden, 2311 EZ, The Netherlands. Her research interests include negotiation, teamwork, and the dynamics of individual agents and organizations. Jonker received her Ph.D. degree from Utrecht University, Utrecht, The Netherlands. She is a fellow and a board member of the European Association for Artificial Intelligence. Contact her at c.m.jonker@tudelft.nl.

M. BIRNA VAN RIEMSDIJK is an associate professor of intimate computing at the University of Twente, Enschede, 7522 NB, The Netherlands. Her research interests include how to take into account our human vulnerability in the design and run-time reasoning of intimate technologies. Van Riemsdijk received her Ph.D. degree in AI from Utrecht University, Utrecht, The Netherlands. Contact her at m.b.vanriemsdijk@utwente.nl.