ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

The Social Context of Human–Robot Interactions

Sydney Thompson,* Kate Candon,* and Marynel Vázquez

Department of Computer Science, Yale University, New Haven, Connecticut, USA; email: sydney.thompson@yale.edu, kate.candon@yale.edu, marynel.vazquez@yale.edu

Annu. Rev. Control Robot. Auton. Syst. 2026. 9:9.1–9.25

The *Annual Review of Control, Robotics, and Autonomous Systems* is online at control.annualreviews.org

https://doi.org/10.1146/annurev-control-030623-015506

*These authors contributed equally to this article

Keywords

context-aware robotics, socially contextual information, social situations, interaction scenarios

Abstract

The human–robot interaction (HRI) community often highlights the social context of an interaction as a key consideration when designing, implementing, and evaluating robot behavior. Unfortunately, researchers use the term "social context" in varied ways. This can lead to miscommunication, making it challenging to draw connections between related work on understanding and modeling the social contexts of human–robot interactions. To address this gap, we survey the HRI literature for existing definitions and uses of the term "social context." Then, we propose a conceptual model for describing the social context of a human–robot interaction. We apply this model to existing work, and we discuss a range of attributes of social contexts that can help researchers plan for interactions, develop behavior models for robots, and gain insights after interactions have taken place. We conclude with a discussion of open research questions in relation to understanding and modeling the social contexts of human–robot interactions.

1. INTRODUCTION

The social context of human–robot interactions is key for the design, evaluation, and automatic generation of appropriate robot behavior (1, 2). It is generally accepted that the social context shapes how humans interpret signals (3), expect others (including robots) to act (4), and behave themselves (5). In the field of human–robot interaction (HRI), prior work has typically investigated some element of robots' understanding of the social context (e.g., 6), how robots should interact or generate suitable behaviors within a particular social context (e.g., 7, 8), and/or human experience with a robot acting within a social context (e.g., 9). Further, academic venues have long invited work that examines interactions in their social context (10), calling for work that is ecologically valid (11) and addresses complex situated encounters (12).

Unfortunately, the use of the term "social context" in HRI is often overloaded or underspecified. For example, "social context" may refer to specific circumstances in which an interaction takes place or a particular application domain (such as healthcare or home robotics). To further complicate the matter, "social context" is often shortened to just "context," which is even more overloaded in practice. The term "context" can refer to information passed to an algorithm to interpret sensor observations [e.g., a map of an environment (13)] or to generate robot behavior [e.g., interaction states provide context to generate locomotion (14)]. Papers in the HRI literature commonly use the term "social context" without describing what it means. Readers experienced in HRI can perhaps make an informed guess, but for others, like new HRI researchers, the meaning of the term could get lost in translation, becoming jargon rather than clarifying the key ideas. The lack of precision makes it difficult to identify, relate, and synthesize gaps in existing work about social contexts and their effects on human–robot interactions. What should we investigate next about the social context of a human–robot interaction? What information about the social context should a robot use for reasoning when deployed in a new scenario? Without having a clear operational definition of what constitutes the social context of a human–robot interaction, it is unlikely that we will be able to answer these types of questions effectively.

Luckily, other related fields have grappled with the problem of conceptualizing the context of human–human and human–computer interactions. For instance, social psychology has long studied the concept of social context when investigating individual human behavior and human–human interactions. Stangor et al. (15) laid the foundation for the scientific study of how social factors influence individual and group behavior. Lewin (16) posited that behavior is a function of individual characteristics and the surrounding social environment, $B = f(P, E)$ (where B is behavior, P is the person, and E is the environment), emphasizing the dynamic interplay between people and their social environments. Work by Argyle et al. (17) has conceptualized social contexts for interactions in terms of social situations. Social situations are characterized by the presence of two or more people who interact and are shaped by a variety of factors, including the roles of the participants, the norms that govern their behavior, their goals or motives, and their forms of communication. As an example of the value of such conceptualization, social situations have recently been adapted to HRI, helping create more varied simulation environments for research on social robot navigation (18).

In the field of human–computer interaction, Dey (19) provided an operational definition of context for ubiquitous computing: "Context is any information that can be used to characterise the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves" (p. 5). This definition served to scope context: For an indoor tour guide application on a mobile device, is the local weather context? According to Dey's definition, information about the weather is not context because the application is being used indoors and the weather does not affect it.

However, information about the presence of other people who take part in the tour with the user is context, because they could affect which sites the user visits while using the application. Based on this notion of context, Dey went on to characterize context-aware computing, helping application builders more easily determine what features their applications should support and what context is critical to support the features. There is similar practical value in better conceptualizing social context in the HRI field.

Riek & Robinson (20) introduced an initial conceptual model of social context for researchers "concerned with the automatic analysis of (and response to) human behavior" (p. 146). They defined social context as "the environment, *E*, where a person, *P*, is situated, with four factors that may influence *P*'s behaviors. These factors include the situational context, *P*'s current social role in *E*, the cultural conventions of both *E* and *P*, and the social norms of *E*" (p. 146). O'Connor & Riek (21) expanded on this definition to provide formalisms for this social context and tease apart the conceptual ideas in practice, and Nigam & Riek (22) applied this model to robotics. Our proposed model builds upon these initial ideas, while focusing specifically on the social context of human–robot interactions.

In this survey, we first review existing definitions and the use of the term "social context" in HRI. The goal is to contrast different perspectives on how the community conceives of social context. We then propose an explicit definition of and conceptual model for the social context of a human–robot interaction. Our conceptual model is inspired by the perspectives about context and social context described previously but is designed specifically for HRI, considering the social nature of human–robot interactions, the importance of relationships within these interactions, and the unique characteristics of robots that differentiate them from other computing technology. To explain our conceptual model, we provide examples of how it can be applied to prior work in HRI and provide a taxonomy for socially contextual information that highlights the diversity of factors it can include.

Overall, with this review and our proposed conceptual model, we provide a pathway to think in an explicit and structured manner about social contexts in HRI. We end by discussing practical uses of our conceptual model and open questions.

2. REVIEW METHODOLOGY

We performed a systematic literature review with two main goals in mind: (*a*) to better understand the uses of the term "social context" in the HRI literature (Section 3) and (*b*) to define and validate our proposed conceptual model for the social context of a human–robot interaction (Section 4). We reviewed publications in well-established venues for HRI work from the years 2012–2023: the proceedings of the ACM/IEEE International Conference on Human–Robot Interaction (HRI), the IEEE International Conference on Robotics and Automation (ICRA), the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), and the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), along with the journal *ACM Transactions on Human-Robot Interaction* (THRI).

The selected venues comprised 27,843 articles, as shown in **Figure 1*a***. We filtered the set of papers for work that discussed aspects of social context in human–robot interactions using string matching. We searched for text matching the regular expression `social(?:ly)?[' ']?context(?:ual)?`. The result was a corpus of 320 papers, some of which were short contribution papers because some conference proceedings combine short and full papers into the same volume. **Figure 1*b*** shows the number of papers in this corpus, organized by year. There is a clear increase in the use of the term "social context" and its close variants (per the regular expression), suggesting that the concept of social context is becoming increasingly important in HRI.

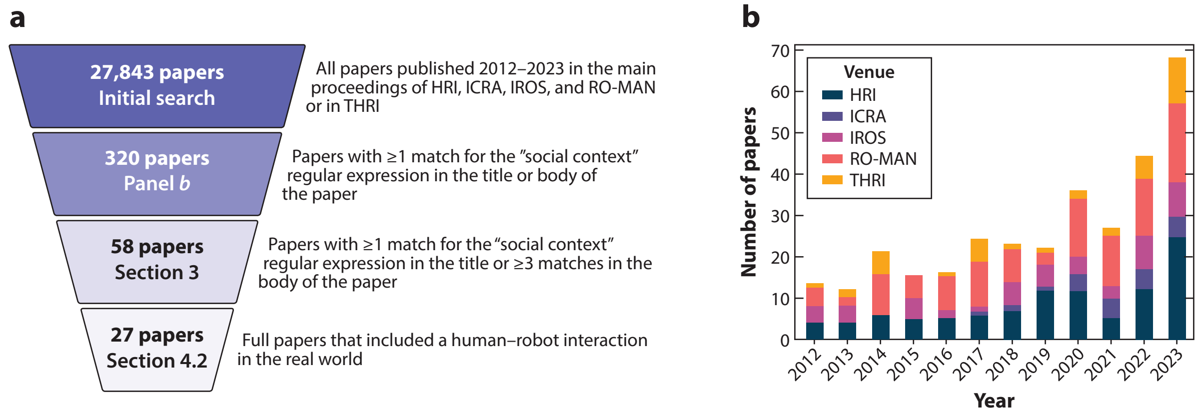

Figure 1

(*a*) Paper selection process for the review. (*b*) Histogram of 320 papers that used the term "social context" (or a close variant, such as "socially contextual") in the title or body of the paper. Abbreviations: HRI, ACM/IEEE International Conference on Human–Robot Interaction; ICRA, IEEE International Conference on Robotics and Automation; IROS, IEEE/RSJ International Conference on Intelligent Robots and Systems; RO-MAN, IEEE International Conference on Robot and Human Interactive Communication; THRI, *ACM Transactions on Human-Robot Interaction*.

The 320 papers included many papers that used the term "social context" without it being a particular focus of the paper; thus, we identified papers that had at least one match to the regular expression for "social context" in the title or at least three matches in the body. This resulted in 58 papers, in which 257 sentences included the term "social context" in the main content, excluding titles and abstracts. In the next section, we analyze these 58 papers to understand how the robotics community uses the term "social context." Later, we narrow down this collection to papers that focus on real-world interactions to explain our proposed conceptual model for social context.

3. THE TERM "SOCIAL CONTEXT" IN HUMAN–ROBOT INTERACTION

This section first discusses different explicit definitions that have been provided for the term "social context" among the 58 papers considered in our review that had at least one match to the regular expression for "social context" in the title or at least three matches in the body. It then discusses broader usage of the term in these 58 papers.

3.1. Explicit Definitions of the Term "Social Context"

In the group of 58 papers, only 4 explicitly defined "social context":

- Huang & Mutlu (23, p. 84) stated that the "social context or setting might characterize the physical environment (e.g., a domestic environment or a workplace), the organization of the interaction (e.g., dyadic interaction or group setting), the relative statuses of the participants (e.g., a supervisor or a subordinate), and the roles of the participants (e.g., a speaker or a bystander)."
- Nigam & Riek (22, p. 3622) defined "the social context for an agent (robot), *P*, in a given environment, *E*, as the disjoint union of several subsets: the situational context as a function of *E*, the social role of *P* in *E*, *P*'s cultural norms (irrespective of *E*)[,]*E*'s cultural norms (irrespective of agents in *P*)[,] and the social norms for *P* in *E*."

- Zaech et al. (24, p. 8988) stated that social context comprises "the positions and velocities of the other agents."
- Luber et al. (25, p. 904) used "the angle of approach α_{ij} between the two [agents] π_i and π_j as the criterion to quantify and distinguish what [they] define here to be a social context."

Huang & Mutlu (23) and Nigam & Riek (22) suggested that the social dynamics between participants and the physical environment play a fundamental role in the social context of an interaction. For social robot navigation, Zaech et al. (24) and Luber et al. (25) proposed that social context comprises the physical behavior of agents, including their relative physical state. Together, the definitions from this subset of the HRI literature highlight three key findings. First, aspects of an environment (e.g., home, work, or the cultural norms of a location) can be an important component of the social context of a human–robot interaction. Second, attributes of individual agents (e.g., their conversational role, position, or velocity) are an important part of the social context. And third, associations between agents (e.g., the relative orientation between two agents or the job hierarchy) can also be part of the social context.

3.2. Broader Usage of the Term "Social Context"

In the corpus of 58 papers, the term "social context" was used in a variety of ways without an explicit definition. Some usages of the term were broad, like using it to describe a problem area or domain (26). Others were more specific, like using it as a synonym for "social norm" (27). Even within the same paper, authors sometimes used the term in different ways, referring to "social context" at different levels of specificity. Because usage varied so widely, it is difficult to suggest a single meaning or implicit definition underpinning the term in the current literature.

To better understand how people currently use the term "social context" in HRI, we systematically examined the uses of the term across the 58 papers. In an initial inspection, we found 6 common usages: state of society, domain, task, social setting, physical setting, and both social and physical setting. We then classified each of the 257 sentences that matched the regular expression for "social context" in the 58 papers into one of these 6 categories. Anything that did not clearly fit into one of these categories was classified as "other." **Table 1** summarizes statistics and examples for each category, which we also describe below:

- State of society: "Social context" refers to the broader state of society or of the world.
- Domain: "Social context" is used to explicitly refer to a broad HRI domain, like healthcare or entertainment. The domain could include many different tasks or environments.
- Task: "Social context" describes a specific application (e.g., robot tutoring), task (e.g., learning to read), or interaction scenario (e.g., a robot asks a child to read something to it).
- Social setting: "Social context" describes the social setting, which might include beliefs, social norms, roles, expectations, or group membership.
- Physical setting: "Social context" describes the physical setting or environment of an interaction, without mentioning social aspects.
- Social and physical setting: "Social context" refers to the social and physical environment of an interaction. In total, 41 papers used the term in this manner across 112 sentences. If both the physical and social setting were referenced explicitly, the paper was further categorized as "social and physical setting (explicit)" (19 papers, 30 total sentences). If the references to the physical or social setting were implicit, then the paper was categorized as "social and physical setting (implicit)" (33 papers, 82 total sentences). Implicit cases include situations

Table 1 Categorization of the usage of the term "social context" in the HRI literature

Category	Stats[a]	Example(s)	References
State of society	4p (6.9%) 7s (2.7%)	"We believe that one key step HRI researchers can take to center the social context is to include a societal implication consideration section in all papers" (28, p. 973).	28–31
Domain	3p (5.2%) 3s (1.2%)	"Robots were previously built to be used in social contexts with members of the public, including healthcare, education, and robots used at home" (32, p. 6882).	26, 32, 33
Task	6p (10.3%) 15s (5.8%)	". . .goals are more important for a specific social context. For instance, if a robot were deployed in a service role that involved interacting with members of the public (e.g., museum tour guide, reception waiter, etc.) . . ." (34, p. 328).	34–39
Social setting	35p (60.3%) 104s (40.5%)	"The parallels between being excluded by robots and being excluded by humans . . .suggests that [robot–robot–human interaction] experiences have the potential to form a powerful social context that impacts humans' emotions and behavior" (40, p. 319).	6, 24–27, 29, 32, 35, 38–64
Physical setting	4p (6.9%) 4s (1.6%)	". . .our robot explored three types of social contexts on our college campus: study areas, dining areas, and lobby areas, across both the student center and library" (65, p. 169).	28, 56, 62, 65
Social and physical setting (explicit)	19p (32.8%) 30s (11.7%)	"Both the limited roles of participants and the confines of the experimental environment present quite a different social context from that in which robots are eventually meant to operate" (47, p. 245).	6, 22, 26, 31, 36, 41, 46–50, 52, 53, 66–71
Social and physical setting (implicit)	33p (56.9%) 82s (31.9%)	"The researchers are immersed within the social context they study, while being aware of the mutual influence researcher and participant have on each other and therefore keeping some distance to the people they study" (72, p. 9).	6, 26–30, 33, 34, 36–38, 41, 43, 44, 47, 48, 56–58, 62, 63, 66, 68–78
Other	7p (12.1%) 12s (4.7%)	"Discriminating and following others' gaze direction is an essential component of establishing a common social context and it is pivotal to the ability to infer others' mental states" (53, p. 152). "Support for this hypothesis would indicate that the agency ascribed to the robot is of key importance in determining whether human interactants will heed its protests, whereas lack of support would indicate the importance of other factors, such as social context" (44, p. 1124).	37, 44, 52, 53, 68, 72, 76

Abbreviation: HRI, human–robot interaction.

[a]The number and percentage of both papers (p) and sentences (s) using the term "social context." There are 257 sentences from 58 papers.

where the physical environment was mentioned in nearby sentences but not in the sentence that contained the term "social context."

- Other: The intended meaning of the term "social context" was unclear or did not fit one of the above categories.

The wide range of usages for the term "social context" motivated us to conceptualize a model, specific to human–robot interactions, that could serve to connect different perspectives in the literature.

4. A CONCEPTUAL MODEL OF THE SOCIAL CONTEXT OF A HUMAN–ROBOT INTERACTION

We propose a conceptual model for the social context of a human–robot interaction. Our goal in creating this model was twofold. First, we wanted to provide an explicit definition of social context specifically for human–robot interactions that connects the literature. We demonstrate how to apply our conceptual model to a variety of use-case scenarios and provide taxonomies for different types of attributes of social contexts, drawing upon the literature. Second, we wanted to provide a tool—the conceptual model—to facilitate planning for interactions, generating behavior, and analyzing interactions after they have occurred. We discuss these practical implications in Section 5 along with future research.

The proposed conceptual model is specifically designed and scoped to describe the social context of a human–robot interaction of interest, where the relevant human(s) and robot(s) act as embodied agents that perform actions in an environment, potentially influencing each other and changing the physical state of the world. For the purposes of the proposed model, a robot is embodied. Physical embodiment makes a robot inherently different from other computing technologies, as discussed in the book *Human–Robot Interaction: An Introduction* (1). While Reeves & Nass's *The Media Equation* (79) suggests that people will act in a similar way with technology as with each other, this is not always true for robots (4).

4.1. Terminology

Because our definition of the social context of a human–robot interaction is dependent on an interaction of interest, we first define what we mean by a human–robot interaction.

Definition 1. A human–robot interaction is an exchange between agents, which must include at least one human and at least one robot. At its core, the interaction corresponds to a sequence of actions taken by the agents in a given environment, which are related to the task of the interaction (or the goals that each agent aims to accomplish). The interaction has temporal bounds that define when it begins and ends.

Human–robot interactions can be explicit, as in conversations, robot tutoring settings, and so on, or implicit (80), occurring without the explicit command or awareness of the human(s) involved in the interaction. For example, common implicit interactions in the social robot navigation literature involve having a robot navigate alone near people (e.g., as in 81). The people are not engaged in co-navigation or in an explicit information exchange with the robot, but they still adapt their actions to the robot as needed.

Human–robot interactions can be one-on-one interactions between a robot and a human, or they can be multiparty, involving more agents (82, 83). They can also change in size over time. The task or agent goals in an interaction can be the same for all agents or different. This can result in collaborative interactions (84), mixed-motive interactions (85), or adversarial interactions (86). Also, interactions can have varied lengths, from brief (like accidental encounters in a given physical space) to longer term (like a robot helping a person practice exercising). In general, we are not concerned with defining these aspects of interactions, but let HRI practitioners and researchers decide what human–robot interaction is of interest, including the relevant agents, task, and temporal bounds.

Given a human–robot interaction of interest, we define the social context of the interaction as follows.

Definition 2. The social context of a human–robot interaction is the set of attributes of the relevant agents, of their environment(s), and of their associations that influence the interaction.

Inspired by how Dey (19) defined "context" for context-aware computing and existing uses of the term "social context" in HRI (Section 3), we consider the attributes in the social context of a human–robot interaction to be information that characterizes the relevant agents, environment(s), and associations between them. The main requirement for these attributes to be part of the social context is that they influence the interaction either directly, by having an effect on the sequence of actions of the interactants, or indirectly, by having an effect on the interacting agents that subsequently influences their actions.

The agents in Definition 2 can be a robot, a person, or some other type of agent. The robot(s) and human(s) that take part in the human–robot interaction of interest have attributes that are part of the social context of their interaction. For instance, this may include the embodiment of a robot, or a person's age, attention, and so on. Other agents could include, for example, pets, whose attributes may be particularly relevant in home or assistive robotics applications.

The inclusion of the term "associations" in Definition 2 is based on the importance that previous work has placed on relationships between agents when defining "social context" for an agent in robotics (21, 22), as well as an emphasis on these associations in context-aware computing (87). In general, associations between agents and environments in Definition 2 can be of three types: agent–agent, agent–environment, and environment–environment associations. Because there is important relational information that can affect human–robot interactions, we make these associations first-class entities in our conceptual model for the social context. That is, we give associations the same level of importance as agent and environment entities. Similar to the latter entities, associations can have more than one attribute that is part of the social context of an interaction. For example, an association between a human and a robot could have information about roles (e.g., whether a person serves as a teacher for a robot) and human impressions about the robot (e.g., whether the person thinks the robot is acting competently). Also, an association between a person and an environment could contain information about the medium through which the person experiences that environment (e.g., in person or via a computer interface with a given set of attributes).

Across human–robot interactions, we find it natural to think about the relevant entities of the interaction via graph visualizations. Nodes in the graph can be of two types (agent or environment), and edges can be used to encode associations between the nodes. If a node or an edge appears in the graph, it is because it has relevant attributes that are part of the social context. For example, imagine an emergency scenario in which a firefighter ($\mathtt{H}_1$) teleoperates a robot ($\mathtt{R}_1$) that goes into a disaster area with fire, as in **Figure 2*a*,*b***. The firefighter teleoperates the robot through a computer interface and from a remote location to stay safe. In this scenario, the social context of the interaction between $\mathtt{H}_1$ and $\mathtt{R}_1$ includes attributes from two environments. First, the environment of the robot ($\mathtt{E}_1$), which the firefighter accesses via a teleoperation interface, influences the commands that $\mathtt{H}_1$ sends to the robot. Second, the environment in which the human is located ($\mathtt{E}_2$) affects the human–robot interaction as well by influencing the human (e.g., noise or other environmental distractions can affect the extent to which the human pays attention to the robot).

The attributes that are part of the social context of an interaction can change over time, and the relevant entities may also change. For instance, in the prior example of a firefighter–robot interaction, it could happen that the robot comes close to a victim ($\mathtt{H}_2$), who then becomes part of the interaction, as in **Figure 2*c***. The attributes that characterize this person and that influence the interaction are now part of the social context.

There can also be additional agents that may not be part of the human–robot interaction of interest but have attributes that are part of the interaction's social context. For example, imagine

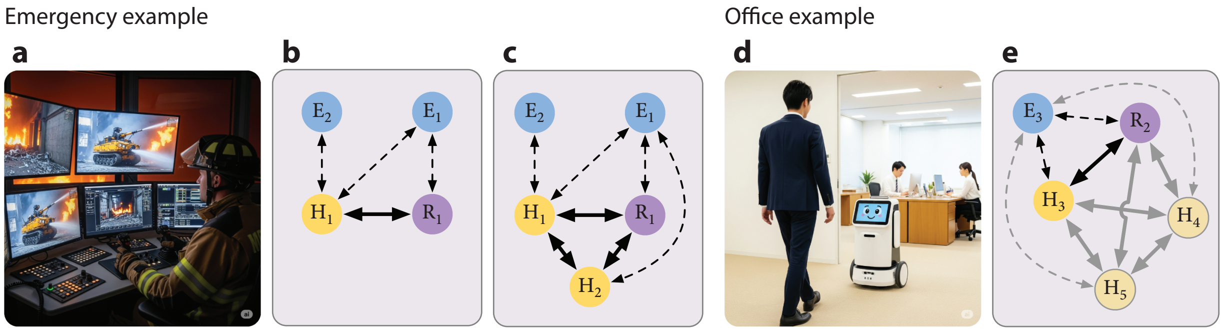

Figure 2

Examples of relevant entities for two different human–robot interactions: (*a–c*) an interaction between a firefighter and a robot and (*d*,*e*) an interaction between a person and a robot guide in an office. Panels *a* and *d* illustrate the interactions at a given time. The diagram in panel *b* shows the relevant entities for panel *a*, including the robot ($\mathtt{R}_1$) and the firefighter ($\mathtt{H}_1$). The firefighter teleoperates the robot that is in $\mathtt{E}_1$ while being physically located in $\mathtt{E}_2$. Dashed arrows indicate relevant agent–environment associations; solid arrows indicate relevant agent–agent associations. The diagram in panel *c* shows the interaction at some time after panel *b*, when the firefighter and robot team find another person ($\mathtt{H}_2$). This illustrates how social context can be dynamic—the relevant elements of the interaction can vary over time, so the attributes that influence human–robot interactions are also variable. In panel *d*, a robot ($\mathtt{R}_2$) brings a person ($\mathtt{H}_3$) to an office, where there are two other people ($\mathtt{H}_4$ and $\mathtt{H}_5$) who are not considered part of the interaction of interest. Hence, they are visualized in panel *e* with a lighter color. However, attributes of $\mathtt{H}_4$ and $\mathtt{H}_5$, like whether they are attending to $\mathtt{R}_2$ or $\mathtt{H}_3$, influence the interaction of interest. As a result, the attributes of $\mathtt{H}_4$ and $\mathtt{H}_5$ are part of the social context of the interaction of interest. Images in panels *a* and *d* generated by Google Gemini 2.0 Flash on May 31, 2025.

that the interaction of interest is between a robot $\mathtt{R}_2$ that guides a person $\mathtt{H}_3$ to a room in an office environment, where other people work (as in **Figure 2*d*,*e***). Then, whether these other nearby people attend to the interaction could influence the actions of $\mathtt{R}_2$ and $\mathtt{H}_3$ by, e.g., making them speak at a lower volume. Thus, the attention of the other people is an attribute of the social context of the interaction of interest.

To more explicitly distinguish the possible attributes of all environments and all agents from those that are part of the social context of a given human–robot interaction, we provide the following definition.

> **Definition 3.** Socially contextual information is the attributes in the social context of a human–robot interaction.

In the remainder of this article, we often describe the social context of a human–robot interaction as a collection, or set, of socially contextual information.

4.2. Selection Criteria for Analysis of Proposed Definition of Social Context

To analyze the current literature through the lens of our proposed definition of social context, we created a subset of the papers selected for the literature review (per Section 2) that focused on aspects of the social context of an interaction. In particular, for each paper in the corpus of 58 papers identified with our matching criteria for the regular expression for "social context," we read the abstract and checked the paper to determine whether it was a full paper about a two-way interaction between at least one robot and at least one human. We did not include videos or extended abstracts in this subset of the papers because these shorter papers are typically about preliminary results. We also excluded papers that discussed only online human–robot interaction studies, where participants did not experience an interaction with a robot in the real world. This filtering process led to 27 papers.

4.3. Example Use-Case Scenarios

This section illustrates how our proposed conceptual model for the social context of a human–robot interaction can be instantiated in specific scenarios. We categorized the 27 papers described in Section 4.2 based on how the paper's main contribution related to our concept of social context. The categories we identified were as follows:[1]

- Study: 10 papers described a study in which socially contextual information constituted independent and/or dependent variables.
- Computational models and systems: 9 papers described a computational model or system that used attributes of the social context of an interaction as inputs or estimated socially contextual information.
- Design: 5 papers discussed the design of human–robot interactions. These discussions included both influences of the social context of a human–robot interaction and how socially contextual information may be influenced.
- Survey: 4 papers were surveys related to ideas captured by our definition of the social context of a human–robot interaction.

For the first three types of contributions described above, we discuss how our model fits the social context discussed in an illustrative paper. While not exhaustive, these use cases demonstrate the flexibility of our conceptual model and its applicability to existing work.

4.3.1. Use case 1: effects of robot appearance. In "Actions Speak Louder than Looks: Does Robot Appearance Affect Human Reactions to Robot Protest and Distress?," Briggs et al. (44) investigated whether a robot's appearance influenced how people responded to the robot verbally protesting a command. In this study, one robot built towers of colored cans. This robot was then removed, and participants were asked to instruct a second robot to knock down the can towers. The demolition robot protested against the participant's request.

In applying our conceptual model to this study, we identified a variety of socially contextual information relevant to the interaction. First, the demolition robot's appearance was the independent variable in the paper, which they found affected the humans' perceptions about the robot's obligation to follow their commands. Thus, under our conceptual model, robot appearance is socially contextual information. Second, we consider the participants' commands to the demolition robot to be socially contextual information. Because the robot executes the participants' commands, the commands directly influence the interaction. Third, the state of the can tower is socially contextual information attributed to the environment. This is because the cans being stacked into a tower impacts the participants' understanding of the task instructions to knock them down. Conversely, the color of the cans is likely not socially contextual information in this case (per Definition 2) because there is no reason to believe that the cans' color affects the interaction.

Under our conceptual model, all positive results in HRI study papers—i.e., all confirmed factors that influence the interaction of interest directly or indirectly—are socially contextual information. These factors are the independent variables of the study and include characteristics of agents (such as action, role, appearance, physical state, or mental state), characteristics of environments (such as location), or associations between agents and/or environments (such as relationships). Negative results indicate that no significant evidence exists that an attribute is socially contextual information.

[1]The numbers per category sum to 28 because one paper fell into both the "study" and "computational models and systems" groups.

4.3.2. Use case 2: selecting listening behaviors. In "A Bayesian Theory of Mind Approach to Nonverbal Communication," Lee et al. (55) introduced a computational model for robot listening behaviors to indicate attentiveness. The robot's behavior is based on the storyteller's actions and a prediction of the storyteller's belief about the robot's attention.

A common assumption in computational models is that the inputs have an underlying causal relationship with the outputs. In this case, there are two models, the outputs of which are the storyteller's beliefs and the robot's action. For the models' inputs where the assumption holds (i.e., the causal relationship exists), the inputs are socially contextual information.[2]

Besides the behavior selection model, the robot's gaze and when it demonstrates a listening behavior are determined by a rule-based model that uses the storyteller's gaze, goals, and attributes of their speech, including pitch, energy, pauses, and length. Because this rule-based model is encoding a causal relationship between the storyteller's behavior and the robot's actions, we consider these input attributes to be socially contextual information.

In addition to the models determining the robot's behavior, Lee et al. (55) discussed specific attributes of the robot that, based on prior research, they thought could influence the interaction, and carefully controlled them to prevent unwanted effects. Specifically, these attributes were the robot's color, gaze, facial expressions, and utterances. These attributes are likely socially contextual information given prior HRI results.

4.3.3. Use case 3: designing interactions for an airport. In "Design Methodology for the UX of HRI: A Field Study of a Commercial Social Robot at an Airport," Tonkin et al. (56) provided a methodology for designing human–robot interactions in public environments that create a positive user experience. They outlined the steps for designing such an interaction and mentioned several factors about environments, humans, and robots that they believed could impact how an interaction would unfold. They noted that, for different deployment locations (e.g., airports, hospitals, or train stations), the volume level can influence whether users can hear any of the robot's speech. For humans, they mentioned that internal state and role (e.g., visitor or staff) can give important insight to their needs, which also may be dependent on the environment. For the robot, they noted that its morphology, personality, task, voice, identity, gestures, and screen display can impact how useful and positive the users' experience is. Under our conceptual model of social context, we would consider each of these attributes likely candidates for socially contextual information.

4.4. Types of Socially Contextual Information

This section presents taxonomies for socially contextual information. We built the taxonomies using the papers described in Section 4.2. For each of the 27 papers, we identified the key attributes of the human(s), robot(s), environment(s), and their associations that were explicitly discussed in the paper. We then used an iterative and collaborative process to develop the taxonomies using affinity diagrams. We expected the 27 papers to provide good coverage for the range of socially contextual information typically considered in HRI, but, when appropriate, we expand with other examples to better convey the richness of the social context of human–robot interactions.

4.4.1. Environment attributes. As shown in **Figure 3**, we identified four main categories of socially contextual information describing the environment of a human–robot interaction:

[2]In practice, validating causal effects for a particular interaction can be difficult, and computational models can be negatively affected by spurious correlations in the data.

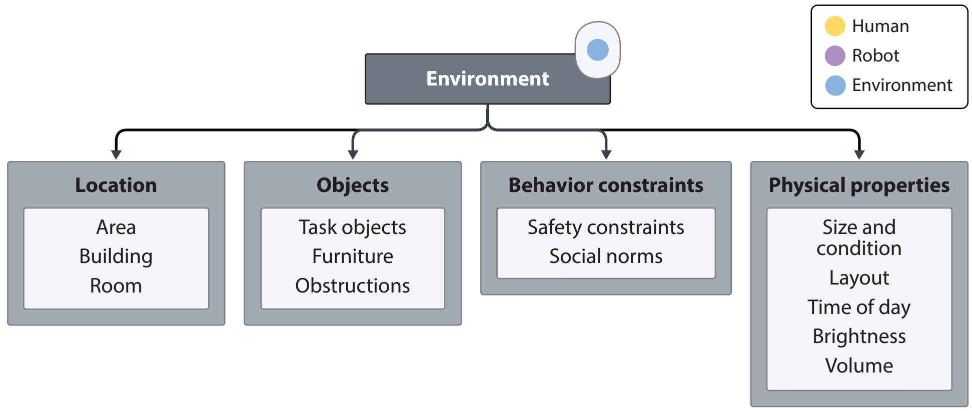

Figure 3

Taxonomy for the socially contextual information of an environment. The blue circle denotes that these attributes could be relevant for the blue environment nodes in **Figure 2**.

- Location: The location was the most common type of socially contextual information for environments. Papers described locations at varying levels of granularity. For example, some described the area in which interactions occur, such as public spaces (56). Other work focused on broader locations by noting the building in which an interaction takes place, such as a nursing home (26) or grocery store (72). Even more specific were references to particular rooms, such as the lobby or activity area within an eldercare facility (71).
- Objects: Humans and robots often engage in physical interactions that involve the manipulation of objects, making attributes of such task objects an important piece of information that can influence the interactions. An example is the location of the cans discussed in Section 4.3.1. Similarly, humans and robots consider furniture (71) or other obstructions (26, 34, 52, 62) when deciding how to navigate within an interaction.
- Behavior constraints: Environments of human–robot interactions can have implications for the behavior of agents, which is socially contextual information that constrains their actions. Two noted examples from our literature review include safety constraints (48), which we view as hard constraints on the behavior of robots, and location-specific social norms (72), which we consider soft constraints.
- Physical properties: We consider the attributes of an environment that can be measured to be their physical properties. When these attributes influence an interaction, they are socially contextual information. For example, papers referenced the layout of an environment as an attribute that can influence behavior. This layout was encoded via environmental maps (62) or referenced through elements such as hallway locations (52). Other physical properties include room size and condition (e.g., whether it needs cleaning) (47), the time of day (22), the brightness of the environment (88), and the volume of the environment (56). Though not directly referenced in the papers in our literature review, this category of environmental attributes could include other properties as well, such as temperature or humidity.

4.4.2. Agent attributes. As shown in **Figure 4**, the papers that we reviewed discussed an agent's actions, common population characteristics, physical state, appearance, and mental states:

- Actions: Action-related attributes can be categorized as either communicative or focused on strategy. Communicative attributes can be associated with verbal and nonverbal actions. The content of utterances (38, 53, 55, 57) and tone of voice (46, 56, 88) were the most common

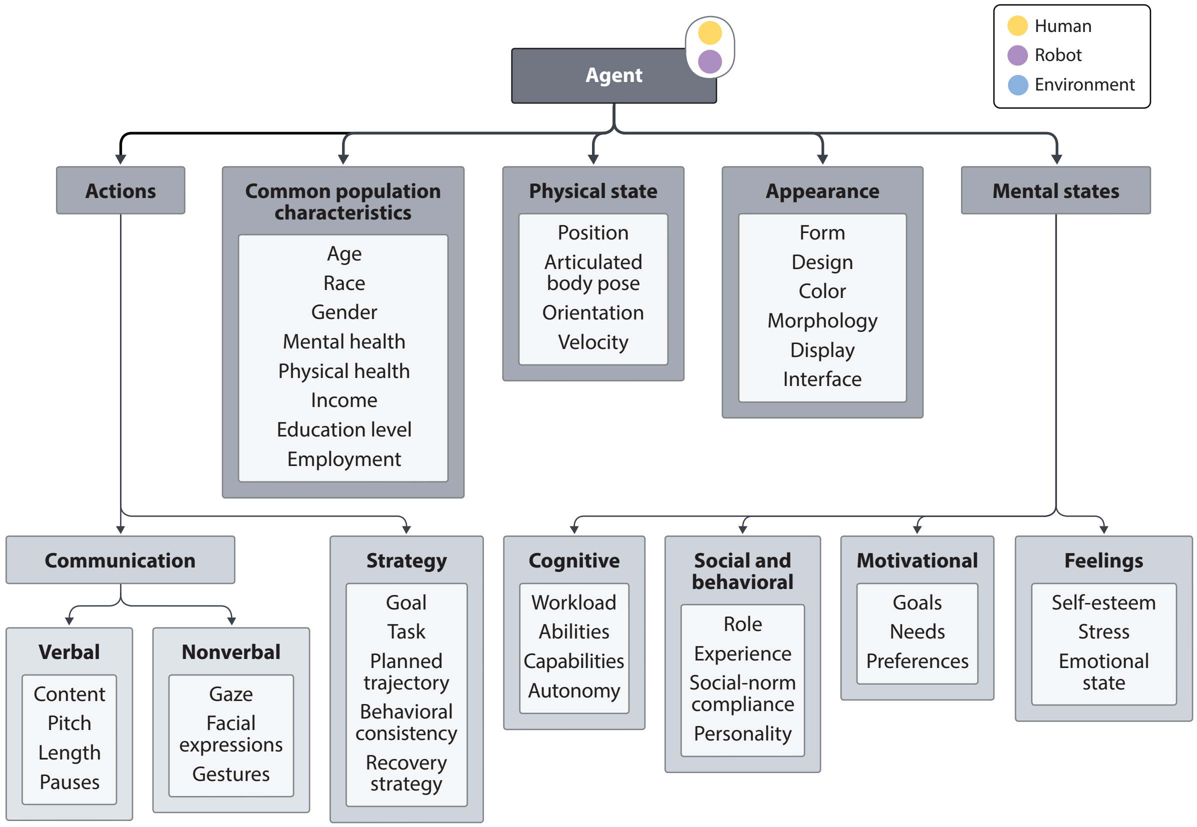

Figure 4

Taxonomy for socially contextual information of an agent. The yellow and purple circles denote that these attributes could be relevant to the human and robot nodes in **Figure 2**.

verbal attributes. Lee et al. (55) also noted how a human's length of utterances or pauses in speech can influence the interaction. Attributes of nonverbal actions commonly include information about gaze (26, 53, 55, 64), facial expressions (55), and gestures (31, 40, 48, 51, 53, 56). Attributes related to an agent's strategy include the agent's goal (6, 34, 62), task (34, 46, 48, 56, 72, 74), planned trajectory (6), behavioral consistency (46), or recovery strategy (46).

- Common population characteristics: A variety of characteristics often used to describe populations—also known as demographics—were often noted for humans in the reviewed papers, although they could be applied to robots as well. For example, the papers that we reviewed noted the age (26, 33, 47, 53, 55, 72, 78), race (71), gender (26, 31, 33, 47, 49, 53, 55, 71, 78), diagnoses related to mental or physical health [e.g., autism (51), dementia (71), other mental health conditions (47), and obesity and diabetes (47)], income (47), education level (47), and employment (47) of humans as factors that could influence interactions. While the selected papers did not explicitly discuss demographic characteristics of robots, it is common for people to assign gender and other similar attributes to robots as they anthropomorphize them. These attributes could influence robots' behavior (e.g., affecting how they communicate with people) or humans' behavior (e.g., affecting mental models of the robots).

- Physical state: While details of the general location of a human–robot interaction can be captured in socially contextual information for the environment, the physical state of an agent is often a piece of socially contextual information. This state often includes an agent's location, which can be represented in (x, y) coordinates (52), but can also have other abstractions. For instance, the physical state of a person could include their articulated body pose (52, 74) or orientation (55, 64). Physical state could also include motion information (52).
- Appearance: An agent's appearance can affect how other agents perceive them. The appearance of robots was often noted in the selected papers. For example, the general design (26, 33, 56), color (55), and morphology (31, 40, 46, 47, 49, 56, 72) of robots can affect interactions. Their screen display (32, 56, 72) or interface (51) can also influence how people interact with them. While the papers that we reviewed did not explicitly discuss the appearance of humans, this attribute can influence human–robot interactions, e.g., in terms of whether a robot can (re)identify a person visually.
- Mental states: Information about the internal, intellectual activity of agents is important socially contextual information because their internal states drive behavior.[3] We identified four subgroups of attributes in the papers that we reviewed. First, mental states can be cognitive factors, such as cognitive workload (31, 46), current abilities to perform specific actions (48, 72), broader capabilities (48), or level of autonomy (47). Second, mental states can be social and behavioral attributes. A common factor is the role of an agent in an interaction. For example, whether a human is a resident, staff member, or visitor at an eldercare facility may influence how a robot chooses to interact with them (56). Similarly, whether a robot is an active or passive participant in a task (e.g., builder versus observer) can influence how a human might interact with said robot (44). Additionally, an agent's experience [e.g., with technology (31, 47) or robots in general (71)], an agent's compliance with social norms (72), or other personality traits (49, 78) can also influence their actions in an interaction. Third, agents have motivational states that can influence their actions, which include attributes related to goals (6, 32, 34, 62), needs (26, 40), and preferences (31). Finally, mental states can relate to an agent's feelings, including self-esteem (40), stress (31), or other emotions (56, 72).

While some specific attributes were discussed more commonly either for robots or for people in the selected papers, these five categories of socially contextual information are generally applicable to both.

4.4.3. Association attributes. As shown in **Figure 5**, attributes of associations generally fall into one of three categories:

- Relationships: Attributes about relationship associations can be of two types—they can encode either information about physical relationships between any entities or information about social relationships between agents. For example, Chang et al. (71) described interactions in different rooms (e.g., a lobby and activity room) in the same building, so the two environments had a physical relationship. Papers also described how the distance between agents or between an agent and an object can influence the agent's actions (62, 64). Attributes related to social relationships include whether or not two agents are partners (57), whether

[3] We considered mental states to be attributed to an agent when they related to the agent itself rather than to other entities. In Section 4.4.3, we discuss mental states that relate to others (e.g., beliefs) as associations.

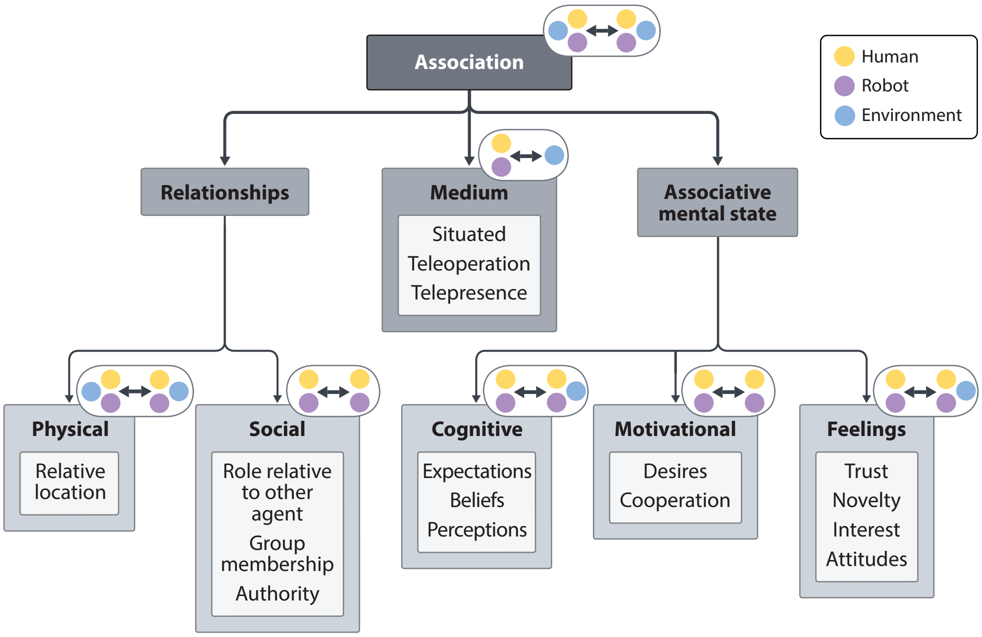

Figure 5

Taxonomy for socially contextual information of an association. The yellow and purple circles denote agent–agent associations, and the blue circles denote associations that involve environments. These colors are consistent with those used in **Figure 2** and in the taxonomies in **Figures 3** and **4**.

two agents belong to the same group (48, 62), and the authority agents have over one another (48).

- Medium: For agent–environment associations, the way in which the agent experiences the environment can influence their behavior and, hence, have an effect on human–robot interactions. For example, it is common for an agent to experience an environment by being situated in it and colocated with other agents (as in **Figure 2*d***). Additionally, agents could experience an environment through computer interfaces, such as a remote teleoperation interface for a robot (as in **Figure 2*a***), a video call, or virtual reality.
- Associative mental state: A variety of internal mental states involve more than one entity and can influence interactions. In the papers that we reviewed, we identified associative mental states that are cognitive factors, such as expectations (46, 48, 72), beliefs (34), or perceptions (44) about other agents or environments. Also, we identified motivational mental states, such as whether a human wants to interrupt a robot's task (72) or touch a robot (48) or thoughts about whether two agents are cooperating as opposed to competing (57). Associative mental states can also be about feelings, such as trust in another agent (46, 48), novelty (72), interest (72), or attitudes (40, 49).

5. DISCUSSION

Given the varied usage of the term "social context" in HRI, we have proposed a conceptual model for the social context of human–robot interactions that can bridge different perspectives. Now, we discuss ways in which we foresee HRI practitioners and researchers leveraging this model in the future and highlight open challenges and interesting future research directions.

5.1. Planning for Human–Robot Interactions

The taxonomies provided in Section 4.4 can serve as an initial checklist for thinking about different kinds of socially contextual information that could influence a human–robot interaction of interest. First, outside of robotics, research has shown that interaction designers can have blind spots to novel conditions that come up during the deployment phase of a technology (89), requiring contingency planning in the interaction design process. Likewise, this could happen in HRI, where the situations that come up during a human–robot interaction can be novel and hard to predict (90, 91). Second, there can be unanticipated human behavior around the adoption of new technology. For instance, while a growing body of research suggests that robots can help support educational efforts (92), research has also indicated that teachers can have negative attitudes toward education robots (93). Attitudes toward robots can be socially contextual information, as described by the taxonomy in **Figure 5**, making it important to plan ahead for them (e.g., by working with stakeholders to facilitate the introduction of robotics technology). Overall, by thinking about a variety of potentially relevant socially contextual information ahead of an interaction based on the proposed conceptual model, HRI practitioners and researchers can prepare for novel situations and challenges that may come up in practice during human–robot interactions.

One difficulty in planning for human–robot interactions is dealing with unknown socially contextual information. Prior work has explored understanding the social contexts of certain environments (94) and developing systems that consider the social context during planning (95, 96). Our taxonomies of socially contextual information are not exhaustive, as there can be more attributes that matter for a given interaction of interest than those reported in the papers that we reviewed. Further, some attributes in our taxonomies may not be relevant to all interactions, and people may change over time, inducing changes in the factors that influence their interactions with robots. For these reasons, we advocate for iterative interaction design processes (e.g., as in 97–99), where interactions are repeatedly prototyped, tested, and refined. Additionally, because user testing can be slow and expensive, it is important for the HRI community to continue innovating in design and evaluation methodologies, which can accelerate understanding of social contexts. These methodologies may include the use of simulations (18, 100), virtual and augmented reality technologies (101, 102), front-end human–robot interfaces (103), crowdsourcing (104, 105), and so on. While these methodologies may not fully replicate real-world results (106–108), they can accelerate the identification of socially contextual information and, hence, help develop better interaction paradigms and robust robotics technology.

5.2. Robot Behavior Generation During Interactions

The proposed conceptual model for the social context of human–robot interactions can aid in developing autonomous robot behavior. Previous work has shown that adapting a robot's behavior to different contexts can improve user experience (7, 109). Intuitively, imagine that a robot had a computational model of the social context of a human–robot interaction—i.e., that it understood what attributes of the relevant agents, environment(s), and their associations influenced the actions of the agents of interest. If the robot could predict the outcome of these influencing effects, it could then generate suitable behavior by searching over its action space for the best actions that help it achieve a desired outcome. This behavior generation setup is generally intractable, but approximations have found value in HRI, e.g., via receding-horizon planning or optimization (110–112).

Computationally representing social contexts and, further, learning the dynamics of interactions in a way that captures all relevant socially contextual information is a difficult challenge. First, a particular interaction may contain extensive amounts of socially contextual information.

Foundational work on computationally modeling the social context of an agent in robotics utilized symbolic representations (21, 22). While effective at describing varied contexts, such representations can be potentially prohibitive from a space perspective due to the explicit nature of the symbolic abstractions. Given recent advances in representation learning with neural networks, there is an opportunity for implicit representations of social contexts to be more scalable. However, because HRI data are scarce, effective generalization will likely require the use of inductive learning biases (113). In particular, we hypothesize that utilizing relational abstractions or graphs (as in **Figure 2**) and machine learning models designed specifically to reason about these abstractions, like graph neural networks (114), could be beneficial for HRI given the importance of associations between agents and environments in the social context of an interaction. Indeed, recent work in HRI has utilized graphs as abstractions to encode attributes of the environment and team dynamics (115), for learning cost functions for motion planning (116), and as state abstractions for learned social robot behavior policies (117).

Second, it is not clear what the best level of specificity is for computationally abstracting socially contextual information. For instance, consider an utterance by a person. It could be computationally abstracted as a high-level intent, text, or a sound wave. Which abstraction is more useful in practice depends on what the robot is trying to achieve during an interaction. For example, high-level intent could be useful for coordinating the robot's behavior with the user (118). The sound wave could aid in synchronizing the robot's speech with the user to build rapport (119). Thus, we suspect that general and efficient computational abstractions for social contexts will ultimately need to be hierarchical. In alignment with this hypothesis, the spatial reasoning community in robotics has advocated for hierarchical, metric–semantic environment maps to enable complex physical robot behavior (120).

Third, some important socially contextual information, such as internal mental states, cannot be directly observed by robots. While machine learning techniques are fueling a variety of approaches for inferring internal human states, e.g., from affective states (121–123) to perceptions of robot behavior (124–126), it remains difficult to measure these internal states in a scalable manner. This poses challenges for building datasets on which to train models and evaluating prediction performance in practice.

Given these challenges and recent advancements in generative AI, it may seem natural to resort to large neural network models, like vision–language models, to generate socially contextual robot behavior. For robot manipulation and navigation, vision–language models built on large-scale internet data are serving as effective backbones for improved scene understanding and are bridging the gap between high-level instructions and low-level control (e.g., see 127, 128). In HRI, large models are becoming increasingly popular to create models of humans (129), generate more varied robot speech (130, 131), and implement a variety of functionality in the control system of a robot (132). However, there are limited data capturing subjective human feedback that can be used in HRI to learn end-to-end robot policies with large models. Additionally, though large models seem to be continually improving, their reasoning remains opaque, making it difficult to understand why, or even when, they make mistakes. This motivates incremental learning approaches focused on continued improvement of robot autonomy (133) as well as utilizing a variety of human feedback (134–136) to adapt or steer robot behavior policies.

5.3. Post-Interaction Analyses

The proposed conceptual model for the social context of human–robot interactions can help to understand interactions after they have taken place. For example, one could analyze the appropriateness of robot behavior based on the relevant social norms that apply to them, which are part of the social context. This idea is in line with prior work in social robot navigation, which has categorized social situations to identify types of interactions where robot performance needs improvement (18), investigated how organizational factors affect the way people respond to delivery robots (137), and classified the environment to adapt robot behavior to different social norms (6). Furthermore, our taxonomies for socially contextual information can help HRI researchers think about potential confounding factors that could lead to incorrect conclusions in experimental HRI work because, by definition, the attributes of the social context influence interactions. Finally, our conceptual model could also aid in defining the concept of generalization in HRI. Research in HRI has called for building a generalizable theory of HRI from in-the-wild social encounters that can provide "a principled understanding of what to expect with different types of robots, performing different types of tasks, in different types of social situations and cultures" (138, p. 2). Our conceptual model for the social context of human–robot interactions can serve to establish a broadly applicable notion of generalization for such theories, where effective generalization entails predicting accurate outcomes in novel social contexts. These novel contexts are characterized by novel values for known socially contextual information, as well as by completely novel environment, agent, and association attributes that influence an interaction of interest.

It would be transformative if robots could reason about causal effects in social encounters based on their observations of interactions and new data that they collect. Recent work has demonstrated the feasibility of learning causal relationships within the HRI domain given a known set of relevant features (139–141), but more work is needed in this direction to capture a wider range of socially contextual information. Ultimately, causal competency, including understanding the social context of human–robot interactions, may be necessary for autonomous, social robots to behave ethically in novel situations (142).

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We are thankful to Nathan Tsoi for helping us create the collection of papers considered in this review. We are also grateful to Anind Dey, Laurel Riek, Bilge Mutlu, Sasha Lew, Houston Claure, and Ulas Berk Karli for their feedback on preliminary drafts. This work was partially supported by the National Science Foundation (grants IIS-2143109 and IIS-2106690) and the US Air Force Office of Scientific Research under the Young Investigator Program (award FA9550-24-1-0085).

LITERATURE CITED

1. Bartneck C, Belpaeme T, Eyssel F, Kanda T, Keijsers M, Šabanović S. 2020. *Human–Robot Interaction: An Introduction*. Cambridge University Press
2. Thomaz A, Hoffman G, Cakmak M. 2016. Computational human-robot interaction. *Found. Trends Robot.* 4(2–3):105–223
3. Vinciarelli A, Pantic M, Bourlard H. 2009. Social signal processing: survey of an emerging domain. *Image Vis. Comput.* 27(12):1743–59
4. Erel H, Vázquez M, Sebo S, Salomons N, Gillet S, Scassellati B. 2024. RoSI: a model for predicting robot social influence. *J. Hum.-Robot Interact.* 13(2):18
5. Aronson E, Wilson TD, Akert RM, Sommers S. 2016. *Social Psychology*. Pearson. 9th ed.
6. Banisetty SB, Rajamohan V, Vega F, Feil-Seifer D. 2021. A deep learning approach to multi-context socially-aware navigation. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE

7. Mutlu B, Osman S, Forlizzi J, Hodgins J, Kiesler S. 2006. Task structure and user attributes as elements of human-robot interaction design. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
8. Kim MJ, Pertsch K, Karamcheti S, Xiao T, Balakrishna A, et al. 2025. OpenVLA: an open-source vision-language-action model. In *Proceedings of the 8th Conference on Robot Learning*, ed. P Agrawal, O Kroemer, W Burgard. Proc. Mach. Learn. Res. 270. PMLR
9. Kim Y, Mutlu B. 2014. How social distance shapes human–robot interaction. *Int. J. Hum.-Comput. Stud.* 72(12):783–95
10. Kiesler S, Hinds P. 2004. Introduction to this special issue on human-robot interaction. *Hum.–Comput. Interact.* 19(1–2):1–8
11. Hoffman G, Zhao X. 2020. A primer for conducting experiments in human–robot interaction. *J. Hum.-Robot Interact.* 10(1):6
12. Bohus D, Horvitz E. 2019. Situated interaction. In *The Handbook of Multimodal-Multisensor Interfaces*, Vol. 3: *Language Processing, Software, Commercialization, and Emerging Directions*, ed. S Oviatt, B Schuller. Morgan & Claypool
13. Vázquez M, Lew A, Gorevoy E, Connolly J. 2022. Pose generation for social robots in conversational group formations. *Front. Robot. AI* 8:703807
14. Liu P, Glas DF, Kanda T, Ishiguro H. 2016. Data-driven HRI: learning social behaviors by example from human–human interaction. *IEEE Trans. Robot.* 32(4):988–1008
15. Stangor C, Jhangiani R, Tarry H. 2014. *Principles of Social Psychology*. BCcampus. 1st. Int. H5P Ed. **https://opentextbc.ca/socialpsychology**
16. Lewin K. 2013. *Principles of Topological Psychology*. Read Books
17. Argyle M, Furnham A, Graham JA. 1981. *Social Situations*. Cambridge University Press
18. Tsoi N, Xiang A, Yu P, Sohn SS, Schwartz G, et al. 2022. SEAN 2.0: formalizing and generating social situations for robot navigation. *IEEE Robot. Autom. Lett.* 7(4):11047–54
19. Dey AK. 2001. Understanding and using context. *Pers. Ubiquitous Comput.* 5:4–7
20. Riek LD, Robinson P. 2011. Challenges and opportunities in building socially intelligent machines. *IEEE Signal Process. Mag.* 28(3):146–49
21. O'Connor MF, Riek LD. 2015. Detecting social context: a method for social event classification using naturalistic multimodal data. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. IEEE
22. Nigam A, Riek LD. 2015. Social context perception for mobile robots. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
23. Huang CM, Mutlu B. 2013. The repertoire of robot behavior: enabling robots to achieve interaction goals through social behavior. *J. Hum.-Robot Interact.* 2(2):80–102
24. Zaech JN, Dai D, Liniger A, Gool LV. 2020. Action sequence predictions of vehicles in urban environments using map and social context. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
25. Luber M, Spinello L, Silva J, Arras KO. 2012. Socially-aware robot navigation: a learning approach. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE
26. Chang WL, Šabanović S. 2015. Interaction expands function: social shaping of the therapeutic robot PARO in a nursing home. In *Proceedings of the 2015 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
27. Kim S, Choi J, Lim Y, Kwak SS. 2022. Should a robot follow social norms? Human-robot interaction design for social relations in mixed age group. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
28. Ostrowski AK, Walker R, Das M, Yang M, Breazeal C, et al. 2022. Ethics, equity, & justice in human-robot interaction: a review and future directions. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
29. Dobrosovestnova A, de Pagter J, Weiss A. 2023. Borrowing, poking and entangling. In search of shared spaces between science and technology studies and human-robot interaction. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM

30. Winkle K, McMillan D, Arnelid M, Harrison K, Balaam M, et al. 2023. Feminist human-robot interaction: disentangling power, principles and practice for better, more ethical HRI. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
31. Šabanović S, Reeder SM, Kechavarzi B. 2014. Designing robots in the wild: in situ prototype evaluation for a break management robot. *J. Hum.-Robot Interact.* 3(1):70–88
32. Alhafnawi M, Hunt ER, Lemaignan S, O'Dowd P, Hauert S. 2022. MOSAIX: a swarm of robot tiles for social human-swarm interaction. In *2022 International Conference on Robotics and Automation (ICRA)*. IEEE
33. Pedersen BKMK, Weigelin BC, Larsen JC, Nielsen J. 2021. Using educational robotics to foster girls' interest in stem: a systematic review. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
34. Briggs G, Scheutz M. 2017. Strategies and mechanisms to enable dialogue agents to respond appropriately to indirect speech acts. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
35. Tuttosi P, Hughson E, Matsufuji A, Zhang C, Lim A. 2023. Read the room: adapting a robot's voice to ambient and social contexts. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
36. Coyne AK, Sapkota K, McGinn C. 2023. "Who said that?" Applying the situation awareness global assessment technique to social telepresence. *ACM Trans. Hum.-Robot Interact.* 12(4):48
37. Duffy BR, Zawieska K. 2012. Suspension of disbelief in social robotics. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. IEEE
38. Van Otterdijk M, Saplacan D, Baselizadeh A, Laeng B, Torresen J. 2023. To shake or not to shake: intuitive reactions of senior adults to a robot handshake in a western culture. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication*. IEEE
39. Bono A, Augello A, Gentile M, Gaglio S. 2020. Social practices based characters in a robotic storytelling system. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
40. Erel H, Cohen Y, Shafrir K, Levy SD, Vidra ID, et al. 2021. Excluded by robots: Can robot-robot-human interaction lead to ostracism? In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
41. Nanavati A, Alves-Oliveira P, Schrenk T, Gordon EK, Cakmak M, Srinivasa SS. 2023. Unintended failures of robot-assisted feeding in social contexts. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
42. Nanavati A, Alves-Oliveira P, Schrenk T, Gordon EK, Cakmak M, Srinivasa SS. 2023. Design principles for robot-assisted feeding in social contexts. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
43. Tagne G, Henaff P, Gregori N. 2016. Measurement and analysis of physical parameters of the handshake between two persons according to simple social contexts. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
44. Briggs G, Gessell B, Dunlap M, Scheutz M. 2014. Actions speak louder than looks: Does robot appearance affect human reactions to robot protest and distress? In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE
45. Navarro I, Oh J. 2022. Social-PatteRNN: socially-aware trajectory prediction guided by motion patterns. *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
46. Tian L, Oviatt S. 2021. A taxonomy of social errors in human-robot interaction. *ACM Trans. Hum.-Robot Interact.* 10(2):13
47. Lee HR, Šabanović S, Chang WL, Nagata S, Piatt J, et al. 2017. Steps toward participatory design of social robots: mutual learning with older adults with depression. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
48. Arnold T, Scheutz M. 2017. Beyond moral dilemmas: exploring the ethical landscape in HRI. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM

49. Erel H, Carsenti E, Zuckerman O. 2022. A carryover effect in HRI: beyond direct social effects in human-robot interaction. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
50. Liu C, Chen Y, Liu M, Shi BE. 2021. AVGCN: trajectory prediction using graph convolutional networks guided by human attention. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE
51. Hwang J, Kim MG, Cho K. 2013. Developing interaction scenarios of robot-mediated imitation training for children with autism spectrum disorders. In *2013 IEEE RO-MAN*. IEEE
52. Jin P, Ohn-Bar E, Kitani K, Asakawa C. 2019. A-EXP4: online social policy learning for adaptive robot-pedestrian interaction. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
53. Pérez-Osorio J, De Tommaso D, Baykara E, Wykowska A. 2018. Joint action with iCub: a successful adaptation of a paradigm of cognitive neuroscience in HRI. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
54. Chuang YS, Hung HY, Gamborino E, Goh JOS, Huang TR, et al. 2020. Using machine theory of mind to learn agent social network structures from observed interactive behaviors with targets. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
55. Lee JJ, Sha F, Breazeal C. 2019. A Bayesian theory of mind approach to nonverbal communication. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
56. Tonkin M, Vitale J, Herse S, Williams MA, Judge W, Wang X. 2018. Design methodology for the UX of HRI: a field study of a commercial social robot at an airport. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
57. Oliveira R, Arriaga P, Alves-Oliveira P, Correia F, Petisca S, Paiva A. 2018. Friends or foes? Socioemotional support and gaze behaviors in mixed groups of humans and robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
58. Chita-Tegmark M, Lohani M, Scheutz M. 2019. Gender effects in perceptions of robots and humans with varying emotional intelligence. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
59. Grimm D, Schörner P, Dreßler M, Zöllner JM. 2023. Holistic graph-based motion prediction. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE
60. Arnold T, Scheutz M. 2018. Observing robot touch in context: How does touch and attitude affect perceptions of a robot's social qualities? In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
61. Moder M, Pauli J. 2022. Proactive robot movements in a crowd by predicting and considering the social influence. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
62. Xu Y, Chakhachiro T, Kathuria T, Ghaffari M. 2023. SoLo T-DIRL: socially-aware dynamic local planner based on trajectory-ranked deep inverse reinforcement learning. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE
63. Jackson RB. 2019. Toward morally sensitive robotic communication. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
64. Vázquez M, Carter EJ, McDorman B, Forlizzi J, Steinfeld A, Hudson SE. 2017. Towards robot autonomy in group conversations: understanding the effects of body orientation and gaze. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
65. Hayes CJ, O'Connor MF, Riek LD. 2014. Avoiding robot faux pas: using social context to teach robots behavioral propriety. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
66. Wang L, Marrella A, Nardi D. 2019. Investigating user perceptions of HRI in social contexts. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
67. Papadakis P, Spalanzani A, Laugier C. 2013. Social mapping of human-populated environments by implicit function learning. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE

68. Kapoor A, Swamy S, Bachiller P, Manso LJ. 2023. SocNavGym: a reinforcement learning gym for social navigation. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
69. Joshi S, Šabanović S. 2017. A communal perspective on shared robots as social catalysts. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
70. Georgiou NC, Ramnauth R, Adeniran E, Lee M, Selin L, Scassellati B. 2023. Is someone there or is that the TV? Detecting social presence using sound. *ACM Trans. Hum.-Robot Interact.* 12(4):47
71. Chang WL, Šabanovic S, Huber L. 2014. Observational study of naturalistic interactions with the socially assistive robot PARO in a nursing home. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE
72. Nielsen S, Skov MB, Hansen KD, Kaszowska A. 2023. Using user-generated YouTube videos to understand unguided interactions with robots in public places. *ACM Trans. Hum.-Robot Interact.* 12(1):5
73. Foster ME, Stuart-Smith J. 2023. Social robotics meets sociolinguistics: investigating accent bias and social context in HRI. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
74. Campbell J, Yamane K. 2020. Learning whole-body human-robot haptic interaction in social contexts. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE
75. Ostrowski AK, Breazeal C, Park HW. 2021. Long-term co-design guidelines: empowering older adults as co-designers of social robots. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
76. Ostrowski AK, Breazeal C. 2022. Design justice for robot design and policy making. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
77. Hennig S, Chellali R. 2012. Expressive synthetic voices: considerations for human robot interaction. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. IEEE
78. Esterwood C, Essenmacher K, Yang H, Zeng F, Robert LP. 2021. Birds of a feather flock together: But do humans and robots? A meta-analysis of human and robot personality matching. In *2021 30th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
79. Reeves B, Nass C. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People*. Cambridge University Press
80. Ju W. 2022. *The Design of Implicit Interactions*. Springer
81. Mavrogiannis C, Hutchinson AM, Macdonald J, Alves-Oliveira P, Knepper RA. 2019. Effects of distinct robot navigation strategies on human behavior in a crowded environment. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
82. Sebo S, Stoll B, Scassellati B, Jung MF. 2020. Robots in groups and teams: a literature review. *Proc. ACM Hum.-Comput. Interact.* 4(CSCW2):176
83. Gillet S, Vázquez M, Andrist S, Leite I, Sebo S. 2024. Interaction-shaping robotics: robots that influence interactions between other agents. *ACM Trans. Hum.-Robot Interact.* 13(1):12
84. Hoffman G, Breazeal C. 2004. Collaboration in human-robot teams. In *AIAA 1st Intelligent Systems Technical Conference*. AIAA
85. Campos J, Alves-Oliveira P, Paiva A. 2016. Looking for conflict: gaze dynamics in a dyadic mixed-motive game. *Auton. Agents Multi-Agent Syst.* 30(1):112–35
86. Duan J, Wang Q, Pinto L, Kuo CCJ, Nikolaidis S. 2019. Robot learning via human adversarial games. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE
87. Zimmermann A, Lorenz A, Oppermann R. 2007. An operational definition of context. In *Modeling and Using Context: 6th International and Interdisciplinary Conference*, ed. B Kokinov, DC Richardson, TR Roth-Berghofer, L Vieu. Springer
88. Tennent H, Moore D, Jung M, Ju W. 2017. Good vibrations: how consequential sounds affect perception of robotic arms. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE
89. Patterson DJ, Sim SE, Aiyelokun T. 2009. Overcoming blind spots in interaction design: a case study in designing for African AIDS orphan care communities. *Inf. Technol. Int. Dev.* 5(4):75–88

90. Murphy RR, Burke JL. 2008. From remote tool to shared roles. *IEEE Robot. Autom. Mag.* 15(4):39–49
91. Honig S, Oron-Gilad T. 2021. Expect the unexpected: leveraging the human-robot ecosystem to handle unexpected robot failures. *Front. Robot. AI* 8:656385
92. Belpaeme T, Kennedy J, Ramachandran A, Scassellati B, Tanaka F. 2018. Social robots for education: a review. *Sci. Robot.* 3(21):eaat5954
93. Reich-Stiebert N, Eyssel F. 2016. Robots in the classroom: what teachers think about teaching and learning with education robots. *Social Robots: 8th International Conference*, *ICSR 2016*, ed. A Agah, J-J Cabibihan, AM Howard, MA Salichs, H He. Springer
94. Taylor A, Matsumoto S, Riek LD. 2020. Situating robots in the emergency department. In *2020 AAAI Spring Symposium on Applied AI in Healthcare: Safety, Community, and the Environment*. AAAI
95. Taylor AM, Matsumoto S, Xiao W, Riek LD. 2021. Social navigation for mobile robots in the emergency department. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE
96. Suresh A, Taylor A, Riek LD, Martinez S. 2023. Robot navigation in risky, crowded environments: understanding human preferences. *IEEE Robot. Autom. Lett.* 8(9):5632–39
97. Lee MK, Forlizzi J, Rybski PE, Crabbe F, Chung W, et al. 2009. The Snackbot: documenting the design of a robot for long-term human-robot interaction. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction*. ACM
98. Hoffman G, Zuckerman O, Hirschberger G, Luria M, Shani Sherman T. 2015. Design and evaluation of a peripheral robotic conversation companion. In *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction*. IEEE
99. Matheus K, Vázquez M, Scassellati B. 2025. Ommie: the design and development of a social robot for anxiety reduction. *ACM Trans. Hum.-Robot Interact.* 14(2):31
100. Hoffman G, Ju W. 2014. Designing robots with movement in mind. *J. Hum.-Robot Interact.* 3(1):91–122
101. Suzuki R, Karim A, Xia T, Hedayati H, Marquardt N. 2022. Augmented reality and robotics: a survey and taxonomy for AR-enhanced human-robot interaction and robotic interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. ACM
102. Walker M, Phung T, Chakraborti T, Williams T, Szafir D. 2023. Virtual, augmented, and mixed reality for human-robot interaction: a survey and virtual design element taxonomy. *ACM Trans. Hum.-Robot Interact.* 12(4):43
103. Toris R, Kammerl J, Lu DV, Lee J, Jenkins OC, et al. 2015. Robot web tools: efficient messaging for cloud robotics. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
104. Lee WY, Sakashita M, Ricci E, Claure H, Guimbretière F, Jung M. 2021. Interactive vignettes: enabling large-scale interactive HRI research. In *2021 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
105. Tsoi N, Hussein M, Fugikawa O, Zhao J, Vázquez M. 2021. An approach to deploy interactive robotic simulators on the web for HRI experiments: results in social robot navigation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE
106. Li R, van Almkerk M, van Waveren S, Carter E, Leite I. 2019. Comparing human-robot proxemics between virtual reality and the real world. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
107. Tsoi N, Sterneck R, Zhao X, Vázquez M. 2024. Influence of simulation and interactivity on human perceptions of a robot during navigation tasks. *ACM Trans. Hum.-Robot Interact.* 13(4):60
108. Esterwood C, Guan RH, Ye X, Robert LP. 2025. Virtually the same or realistically different? A meta-analysis of real versus "not so real" robots. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
109. Porfirio D, Sauppé A, Albarghouthi A, Mutlu B. 2020. Transforming robot programs based on social context. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM
110. Ng E, Liu Z, Kennedy M. 2023. It takes two: learning to plan for human-robot cooperative carrying. In *2023 IEEE International Conference on Robotics and Automation*. IEEE
111. Ghose D, Lewkowicz M, Dong D, Cheng A, Doan T, et al. 2024. Planning with critical decision points: robots that influence humans to infer their strategy. In *2024 33rd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE

112. Renz H, Krämer M, Bertram T. 2024. Moving horizon planning for human-robot interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
113. Mitchell TM. 1980. *The need for biases in learning generalizations*. Tech. Rep., Department of Computer Science, Rutgers University
114. Battaglia PW, Hamrick JB, Bapst V, Sanchez-Gonzalez A, Zambaldi V, et al. 2018. Relational inductive biases, deep learning, and graph networks. arXiv:1806.01261v3 [cs.LG]
115. Altundas B, Wang Z, Bishop J, Gombolay M. 2022. Learning coordination policies over heterogeneous graphs for human-robot teams via recurrent neural schedule propagation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
116. Manso LJ, Jorvekar RR, Faria DR, Bustos P, Bachiller P. 2021. Graph neural networks for human-aware social navigation. In *Advances in Physical Agents II: Proceedings of the 21st International Workshop of Physical Agents, November 19–20, 2020*, ed. LM Bergasa, M Ocaña, R Barea, E López-Guillén, P Revenga. Springer
117. Gillet S, Thompson S, Leite I, Vázquez M. 2025. Templates and graph neural networks for social robots interacting in small groups of varying sizes. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
118. Huggins M, Alghowinem S, Jeong S, Colon-Hernandez P, Breazeal C, Park HW. 2021. Practical guidelines for intent recognition: BERT with minimal training data evaluated in real-world HRI application. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
119. Nishimura S, Nakamura T, Sato W, Kanbara M, Fujimoto Y, et al. 2021. Vocal synchrony of robots boosts positive affective empathy. *Appl. Sci.* 11(6):2502
120. Hughes N, Chang Y, Hu S, Talak R, Abdulhai R, et al. 2024. Foundations of spatial perception for robotics: hierarchical representations and real-time systems. *Int. J. Robot. Res.* 43(10):1457–505
121. Aly A, Tapus A. 2012. Towards an online fuzzy modeling for human internal states detection. In *2012 12th International Conference on Control Automation Robotics and Vision*. IEEE
122. Staffa M, Rossi S. 2022. Enhancing affective robotics via human internal state monitoring. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
123. Churamani N, Barros P, Gunes H, Wermter S. 2022. Affect-driven learning of robot behaviour for collaborative human-robot interactions. *Front. Robot. AI* 9:717193
124. Cui Y, Zhang Q, Knox B, Allievi A, Stone P, Niekum S. 2021. The EMPATHIC framework for task learning from implicit human feedback. In *Proceedings of the 2020 Conference on Robot Learning*, ed. J Kober, F Ramos, C Tomlin. Proc. Mach. Learn. Res. 155. PMLR
125. Stiber M, Bohus D, Andrist S. 2024. "Uh, *this* one?": leveraging behavioral signals for detecting confusion during physical tasks. In *Proceedings of the 26th International Conference on Multimodal Interaction*. ACM
126. Zhang Q, Tsoi N, Nagib M, Choi B, Tan J, et al. 2025. Predicting human perceptions of robot performance during navigation tasks. *ACM Trans. Hum.-Robot Interact.* 14(3):46
127. Vuong Q, Levine S, Walke HR, Pertsch K, Singh A, et al. 2023. Open X-Embodiment: robotic learning datasets and RT-X models. In *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition @ CoRL2023*. **https://openreview.net/forum?id=zraBtFgxT0**
128. Song D, Liang J, Payandeh A, Raj AH, Xiao X, Manocha D. 2024. VLM-Social-Nav: socially aware robot navigation through scoring using vision-language models. *IEEE Robot. Autom. Lett.* 10(1):508–15
129. Zhang B, Soh H. 2023. Large language models as zero-shot human models for human-robot interaction. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE
130. Skantze G, Irfan B. 2025. Applying general turn-taking models to conversational human-robot interaction. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE
131. Kim CY, Lee CP, Mutlu B. 2024. Understanding large-language model (LLM)-powered human-robot interaction. In *Proceedings of 2024 ACM/IEEE International Conference on Human-Robot Interaction*. ACM
132. Williams T, Matuszek C, Mead R, Depalma N. 2024. Scarecrows in Oz: the use of large language models in HRI. *ACM Trans. Hum.-Robot Interact.* 13(1):1
133. Churamani N, Kalkan S, Gunes H. 2020. Continual learning for affective robotics: Why, what and how? In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE

134. Chetouani M. 2021. Interactive robot learning: an overview. In *Human-Centered Artificial Intelligence: Advanced Lectures*, ed. M Chetouani, V Dignum, P Lukowicz, C Sierra. Springer
135. Chernova S, Thomaz AL. 2022. *Robot Learning from Human Teachers*. Springer
136. Fitzgerald T, Koppol P, Callaghan P, Wong RQJH, Simmons R, et al. 2022. INQUIRE: INteractive Querying for User-aware Informative REasoning. In *Proceedings of the 6th Conference on Robot Learning*, ed. K Liu, D Kulic, J Ichnowski. Proc. Mach. Learn. Res. 205. PMLR
137. Mutlu B, Forlizzi J. 2008. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*. ACM
138. Jung M, Hinds P. 2018. Robots in the wild: a time for more robust theories of human-robot interaction. *ACM Trans. Hum.-Robot Interact.* 7(1):2
139. Edström F, Hellström T, De Luna X. 2023. Robot causal discovery aided by human interaction. In *2023 3nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE
140. Castri L, Mghames S, Hanheide M, Bellotto N. 2022. Causal discovery of dynamic models for predicting human spatial interactions. In *Social Robotics: 14th International Conference, ICSR 2022*, ed. F Cavallo, J-J Cabibihan, L Fiorini, A Sorrentino, H He, et al. Springer
141. Castri L, Mghames S, Hanheide M, Bellotto N. 2023. Enhancing causal discovery from robot sensor data in dynamic scenarios. In *Proceedings of the Second Conference on Causal Learning and Reasoning*, ed. M van der Schaar, C Zhang, D Janzing. Proc. Mach. Learn. Res. 213. PMLR
142. Hellström T. 2021. The relevance of causation in robotics: a review, categorization, and analysis. *Paladyn J. Behav. Robot.* 12(1):238–55