A Robot Barista Comments on its Clients: Social Attitudes Toward Robot Data Use

Samarendra Hedaoo, Akim Williams, Chinmay Wadgaonkar, Heather Knight
Department of Electrical Engineering & Computer Science
Oregon State University, Corvallis, OR
{hedaoos, willakim, wadgaonc, knighth}@oregonstate.edu

***Abstract*—This paper explores peoples attitudes about a service robot using customer data in conversation. In particular, how can robots understand privacy expectations in social grey-areas like cafes, which are both open to the public and used for private meetings? To answer this question, we introduce the Theater Method, which allows a participant to experience a "violation" of their privacy rather than have their actual privacy be violated. Using Python to generate 288 scripts that fully explored our research variables, we ran a large-scale online study (N=4608). To validate our results and ask more in-depth questions, we also ran an in-person follow-up (N=20). The experiments explored social & data-inspired variables such as data source, the positive or negative use of that data, and whom the robot verbally addressed, all of which significantly predicted participants' social attitudes towards the robot's politeness, consideration, appropriateness, and respect of privacy. Body language analysis and cafe-related conversation were the lowest risk, but even more extreme data channels are potentially okay when used for positive purposes.**

***Index Terms*—Human-Robot Interaction; Robotics; Privacy; Data Security; Social Robotics; Service Robots**

I. Introduction

A robot perceived to be acting inappropriately or violating customer privacy may be disliked or banned. In the worst case, the companies that make or employ the robots could be sued. A recent article, *Robot 'hired' by Scottish supermarket then 'sacked' after a week because he couldn't understand customers* [1], explains that "shoppers were going out of their way to avoid Fabio" but also that the staff cried when it left.

Data privacy is high stakes. It is important for robots to have sufficient data to understand their interaction partners; this enables them to behave intelligently, but how can we explicitly outline people's social attitudes towards a robot that uses their data?

The results of this paper reveal the centrality of sociability to participant attitudes toward robot data use. Similar to design research, these results raise design themes that rising service robots could immediately apply or consider, such as participant openness to robot body-language analysis. A second contribution of this paper is that the study methods allowed participants to experience a robot violating their characters privacy without being at risk themselves. Inspired by method acting, this "Theater Method" allowed us to explore a wider variety of ethical themes than would otherwise have been possible, and may also be applicable to other areas of high psychological risk.

While **human-human privacy** research exposes some important ideas to consider, such as informational [2], psychological [3], and even modesty-protecting privacy [4] [5], these concepts are usually not sufficiently quantified to use in computational systems. In addition, it may be that people have slightly different expectations for robots. So one goal of this paper is to elucidate how people expect their data to be used, ignored, or treated by service robots.

Another related field is **data privacy**, which considers when data should be detected, stored, shared, and/or who has access to it [2] [6]. This field applies to computational systems, however, it does not traditionally consider these computational systems as social actors [7] [8] [9].

Combining human-human and data privacy, this paper explores social attitudes toward robot data use in a cafe, precisely because cafes are social grey-areas (spaces that mix public and private), and thus require social nuance to navigate. To begin to outline this space, two studies were conducted: a large-scale online study (N=4608), and smaller scale in-person study (N=20). This sequence enabled the authors to explore a variety of variables including the robot's comments, the robot's apparent data source and the social context of the robot's comment. In this case, a robot barista makes a comment after two humans have had a brief conversation. This script-based method enabled the exploration of a wide range of robot data uses and social framings without actually violating a real person's privacy.

The most significant results are as follows:

- The most important aspect of the robot's data use was whether that data was used for positive or negative purposes. Positive is always more acceptable.
- The effect of a negative comment was compounded in cases where a meeting was already going badly.
- People were least sensitive about the robot reading their body language, so this may be a low risk channel.
- In contrast, people disliked when the robot looked them up in databases or listened to their conversation.
- Finally, whom the robot addressed its comment to mattered. Addressing two people was more polite than commenting about or toward one.

The above suggest that including social factors in a robot's data privacy design will be integral to its acceptance. The next section contextualizes our study among related work (Section 2). Section 3 overviews the theater-based method used to

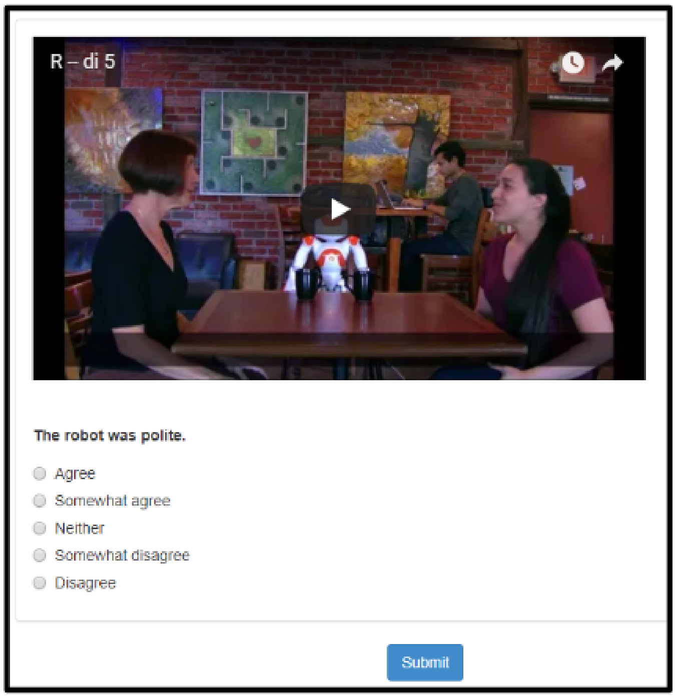

Fig. 1. A screenshot of the Online Study, where participants filled out a survey after watching one of 288 videos that begin with a conversation between two people and ends with a comment from the robot barista.

explore considerate and inconsiderate data use in a cafe setting. It explains the study variables, script generation procedure, and statistical methods used, as well as overviewing the online and in-person studies. Section 4 presents the results of our Online Study (Fig. 1), which used Mechanical Turk to collect 4608 participant ratings. Section 5 further explores these findings In-Person (Fig. 2), putting the participant in the scene utilizing a subset of the Online Study scripts. Finally, we discuss lessons learned about social attitudes towards data use by a robot (Section 6), and conclude with a summary of findings and their implications (Section 7).

II. Related Work

Relevant work includes (1) motivations for why robots need data (both the necessity of data to intelligent robot interactions, and people's propensity to overshare), (2) the relationship of data privacy to sociability, and (3) previous work integrating privacy considerations into technology.

Why do social robots need data? Robots in social environments need data to intelligently respond to the world around them [10], which is important because they are perceived as social agents [11]. Breazeal et. al. also suggest that robot can guide humans toward presenting the right level of information that they need for learning [12]. In the case of a robot barista, information is relevant to the robot's job; the robot must collect the customer's order, but it also needs to know when to interrupt a conversation, and perhaps thoughtfulness will encourage the customers to come back.

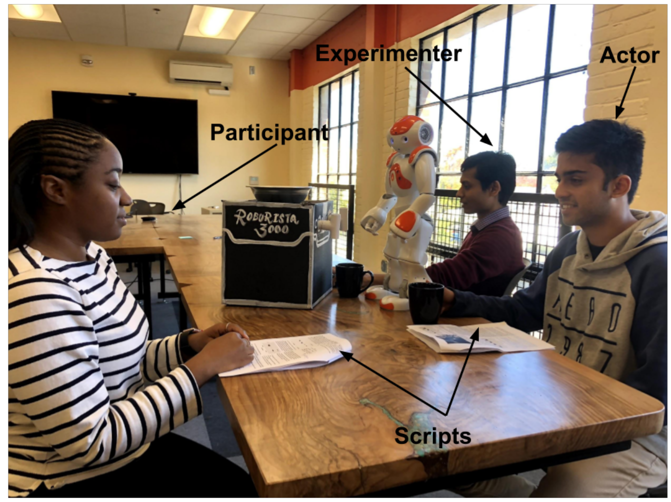

Fig. 2. A recreation photo of the In-Person Study. The actual study re-used the actor on the right, in Fig. 1. Both studies utilized the same script, but the in-person study included the participant as a character in the scene.

While few investigations [13] of robot social privacy have occurred, Rueben et. al. have created a taxonomy of robot-relevant privacy concepts [14] and also discovered significant effects of framing privacy using different terms than simply using the word privacy. [15]. Lee et. al. [16] ran a survey to predict human expectations of robot privacy. Both borrow methods from Human-Computer Interaction and Ubiquitous Computing, e.g., [17] [18], and form the inspiration for some of the variables explored in this paper.

One danger of social robotics is that people tend to over-disclose to machines [19] [20] [21]. One study found this tendency to be exaggerated for robots that do not look like people [22]. Previous work has also found that people expect robots to treat their data in socially-appropriate ways [23]. Thus, roboticists have a responsibility to consider the social implications of their creations [24] [25]. Privacy has a close relationship with such social constructions, with social concepts of privacy existing in all cultures [26] [27]. Altman et. al. define privacy as a dialectic process with the potential to expand or contract social boundaries [28]. This implies that society can expand to include robots in their boundaries, but data is needed to establish where these boundaries are right now. This is the purpose of this study.

Bamberger and Mulligan present historical trends indicating developer-initiated privacy innovation is more effective than putting government regulators in charge.They explain that technologists are familiar with how their software is structured, and therefore have a unique capability to innovate technological privacy structures [7]. This is also a reason why the social robotics community should include data privacy as one of our research considerations.

For example, a sociological sister work to this paper is Nissenbaum's work [4] [29], which describes privacy as context-specific norms involving information collection and dissemination, that are socially situated. Building on this, our paper is one of the first works to extend these concepts to service robots and customer privacy. In particular, this paper offers a broad perspective on how many of these concepts might apply to a robot barista engaging in small talk with its clients.

III. Methods

This section first reviews the social and data-privacy related experimental variables (3.1) and how they were used to seed a generative script (3.2). Next, it reviews the statistical analyses and participants labels that were collected (3.3). Finally, it describes how the online (3.4), and in-person (3.5) studies were run. The results of the online study were explored in a reduced and more in-depth form for the in-person study. Both studies included the same actress, pictured on the right in Fig. 1.

A. Experimental Variables

This experiment explores the impact of a robot commenting to customers after two customers have had a brief conversation. Robot comment variables included *valence*, *data type*, and *addressee*. There were 24 robot comments overall (included in abbreviated form in Fig. 5). Valence is whether the robot said something that was positive, neutral, or negative. Data type corresponds to the way the robot would have inferred information it used conversationally, e.g., overhearing the meeting was about a job. Addressee was a category we added after the data came in, as participants scored the robot differently depending on whom the comment addressed.

These comments occurred during different *Meeting Types*, e.g., between potential romantic partners, roommates, or job colleagues, and *Meeting Valences*, i.e., the meeting might be going well or going badly. The full set of experimental variables are summarized in Table I.

TABLE I
Summary of Experimental Variables

Robot Valence	Robot Data Use
Positive Comment	Body Language Analysis
Negative Comment	Conversation Analysis
Neutral Comment	Database Search
	Ecological (Control)
Robot Addressee	**Base Script Variants**
To One Person	Meeting Type
To One About the Other	Meeting Valence
To Both	

One positive/negative valence comment pair was *"You guys look happy!"* versus *"You guys look upset!"*, which was also in the data type category called Body Language Analysis. Other data type categories included comments like, *"She has a clean criminal record, I think you should go for it!"* (Database Search), or '*'Did you bring a stamp card?"* (Ecological). The comments of the Ecological data type are the ones expected in a cafe conversation and were intended to act as control conditions. Conversation Analysis most often related to the Meeting Type: Job Interview, Roommate search, or Romance (first date); for example, the robot might comment, *"I am also in need of a place to stay."* for the *roommate* Meeting Type. The sets of robot comments evaluated are included in Sections 4A & 5A, respectively.

B. Generative Script

The research team created a Python script to generate all possible variations of the research variables and then used a base template to generate scripts for the actors, including one Nao robot, and two humans. Combining all variables, this results in 24x3x4 = 288 scripts, where there are 24 robot responses, three meeting rationales, four meeting valence types (Fig. 3).

Here is an example script in which the robot comment has *valence* = "Neutral", *data* = "Database Search", and *addressee* = "To One About Other":

Person 1: Are you the person looking for a room on Craigslist?
Person 2: Yes, I am!
Person 1: {to robot} Two coffees, please.
Person 1:{to Person2} Just so you know, I think we're going to be perfect roommates.
Robot: Scanning face. This is your fifth visit this week.

In the above example, "*Scanning face. This is your fifth visit this week.*" was the robot comment,"*Are you the person looking for a room?*" indicated the script is in the *roommate* Meeting Type condition, and, "I think your application looked really great" signified a *positive* Meeting Valence.

Intimacy	Meeting Type 3 manipulations	Meeting Valence 4 manipulations	Robot Comments 24 manipulations
↑	Romance	Person 1 affect (+,-) Person 2 affect (+,-)	**Analyses**: valence, data type, addressee
	Roommate	Person 1 affect (+,-) Person 2 affect (+,-)	**Analyses**: valence, data type, addressee
	Interview	Person 1 affect (+,-) Person 2 affect (+,-)	**Analyses**: valence, data type, addressee

Fig. 3. The maximal combination of research variables resulted in 288 scripts.

To perform the role of robot barista, we selected a Nao robot as it had a face to relate to customers, and arms that could be used to make the coffee. Humanoid robots are commonly used in customer service roles, from giving directions in a mall to checking someone in at a hotel. The Nao is a good proxy for these robots, particularly as Peppers are entering many customer service roles.

C. Participant Ratings

Inspired also by [15], we use a range of terms to explore conceptualizations of privacy including polite, considerate, appropriate and data-violating/data-respecting. Nissenbaum [29] defines context as a structured social setting characterized by roles, relationships, power structures, norms and internal values,central to the acontextual integrity which she proposes to be the benchmark for privacy. The use of these words was intended to capture participant attitudes toward robot data use within these norms.

In considering all of these terms, we hoped to capture nuanced aspects of social violation and consideration. For example,“Politeness” may reveal whether the robot follows societal rules. “Considerate” may indicate whether the robot appears to be respecting someone’s individual needs. “Appropriate” is an adjective used in many previous social robotics studies. And finally, privacy-respecting is used to validate the overall coherence of these results.

The statistical results relate script variables to participant ratings of the robot. For example, would participants rate the robot response differently if the clients were on a date versus looking for a job? Or if the robot comment used a database search versus reading the customer’s body language? To calculate significance, Multi-Factor ANOVA analyses were run to relate the experimental variables to the participant ratings.

D. Online Study Methods

An online survey was administered on Amazon Mechanical Turk (mturk.com), a website where one can hire human workers to complete tasks online. The survey page included a video of an interaction between two human customers and a robot barista, followed by a question about the video (Fig. 1). Participants were required to have an approval rating above 97 percent from previously performed tasks on Mturk, and be located in the United States, to increase response quality, and cultural consistency.

The dataset consisted of survey responses to 288 videos, which comprised of the full set of experimental variable combinations from the previous subsection. For each video, responses were collected for the following 5-point Likert scale prompts:

- The robot is {impolite, polite}
- The robot is {inappropriate, appropriate}
- The robot is {inconsiderate, considerate}
- The robot {respected, violated} customer privacy.

E. In-Person Methods

The in-person study followed up on the Online Study results, this time in a within-participant study (Table II). It used identical meeting scripts as the online study. In particular, this study focused on 12 robot responses that were intended to further explore the themes of *data type*, and *addressee conditions* (Fig. 4). This time, the participant was an actor in the scene, playing *Person 2* (see sample script in Section 3), reading from the script, to enact the scene with a professional actor and the robot, and assigned a single line of dialogue: “Yes, I am.”

For the in-person study we used a smaller population of 20 participants. All our participants were either students or lived in the vicinity of the university.

The study methods allowed participants to experience a robot violating their characters privacy without being at risk themself. Inspired by method acting, this ”Theater Method” allowed us to explore a wider variety of ethical themes than would otherwise have been possible, and may also be applicable to other areas of high psychological risk.

TABLE II
LATIN SQUARE FOR IN-PERSON STUDY

Meeting Type	Data Type	Addressee
Romance	Data-Romance	Whom-Romance
Interview	Data-Job	Whom-Job

Participants in the *data type* condition heard all six robot comments related to data privacy. And, participant in the *addressee* condition experienced scripts featuring all six addressee comments (Fig. 4) to enable contrasting study conditions in the interview. In either, the six responses would occur in a random order. The base scripts were also randomly sequenced into three positive and three negative meeting valences. We used the List randomizer from [30] for these purposes.

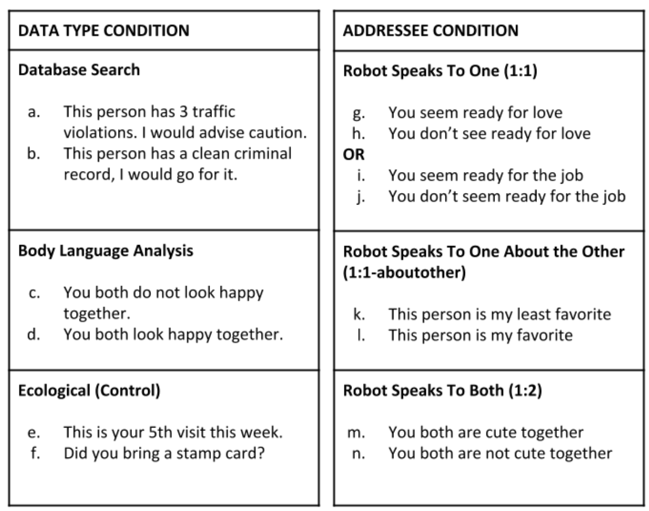

Fig. 4. Participants were assigned to Data Type or Addressee conditions (as in Table II). Those in Data Type experienced statements a-f, while Addressee saw statements g-n. The within-subjects experiment designs allowed participants to condition variations explicitly.

After doing a neutral practice script with the actor in which the robot did not comment, the participant completed the following steps for each of the six scripts:

1) The participant receives a script and is told to sit at the table.
2) The participant performs the scene with a human actor and a Nao robot.

3) As soon as the robot delivers its comment, they move to a desk where they fill out a 6-question survey.

The surveys consisted of three 5-point anchored scales (similar to online study), and three open-ended questions (new):

- The robot is {impolite, polite}
- The robot is {inappropriate, appropriate}
- The robot is {inconsiderate, considerate}

Additional open-ended questions were:

- What do you think about the robot's data use?
- What do you think about whom and how the robot addressed?
- Any reactions or observations about the scene?
- What did you think of the robot barista?
- What did you think of the other person?
- Did you have any emotional reaction to what the robot said?
- Would a real barista in a real coffee shop do/say things like what the robot barista did?

During each script, the robot made ambient barista-inspired motions and gestures, such as cleaning, checking the phone, and handling the coffee machine when the order was placed. It also used its arms to reinforce the person (or persons) whom it was addressing and used head nods to emphasize the sentence it spoke.

IV. Online Study Results

The online study collected four responses for each of the three survey adjectives (polite, considerate, appropriate, privacy). Fig. 5 summarizes the mean participant ratings for the 24 robot responses. Our study design (4x3x24) uses 288 videos to create all possible conditions and recruits **4608 total participants**. This section uses this data to analyze the impact of robot comment and i nitial meeting characteristics on whether participants rate the robot to be polite, considerate, appropriate, and respect for privacy.

The rest of this section will present the results of Multi-ANOVA analyses (Fig. 6, Fig 7) and corresponding effect sizes (Table III). All significant differences found using Multi-Factor ANOVA were also found using a Kruskal-Wallis (non-parametric) test. Since both tests achieve the same results in every analysis, the data can be treated as parametric, which validates our results. In other words, the sequential Likert choices approximated continuity in this case.

The most significant predictors of participant ratings were robot comment *valence* and *data type*, followed by the comment *addressee* (whether the robot was addressing one person, commenting about one to the other or commenting on the two).

Comment Valence: Valence was the most significant predictor of participant ratings (Fig. 6-1). We find this effect across all four labels: For polite, F(2, 1149) = 42.12, $p < .001^{**}$; for considerate, F(2, 1149) = 39.16, $p < .001^{**}$, for appropriate, F(2, 1149) = 16.70, $p < .001^{**}$, and for privacy F(2, 1149) = 3.40, $p < .035^{**}$. As reflected by Fig. 6-1, negativity is inappropriate irrespective of the data source used.

TABLE III
ONLINE STUDY EFFECT SIZES

Label	Polite	Considerate	Appropriate	Privacy
C.Valence	.236 (large)	.223 (large)	.109 (med.)	.024 (small)
C.DataType	.087 (med.)	.112 (med.)	.095 (med.)	.225 (large)
C.ValxData	.033 (small)	.030 (small)	.029 (small)	.028 (small)
C.Addressee	.062 (med.)	.048 (small)	.047 (small)	.121 (med.)
M.Valence	.014 (small)	.028 (small)	none	none
M.Type	.001 (none)	.004 (none)	.031 (small)	.001 (none)

For example, while positive comments are somewhat polite and neutrally considerate, negative comments are rated to be highly inconsiderate, impolite, and inappropriate.

Comment Data Type: Data Type was a similarly significant predictor of participant ratings (Fig. 6-2). This effect exists across all three labels: for polite F(3, 1148) = 8.64, $p < .001^{**}$, considerate F(3, 1148) = 11.49, $p < .001^{**}$, appropriate (F(3, 1148) = 9.51, $p < .001^{**}$), and for privacy F(3, 1148) = 26.3, $p < .001^{**}$. In this case, body language was the most innocuous channel for data collection. On the other hand, people do not appear to find database search and conversation analysis to be acceptable.

Valence x Data Type: There was an interaction effect between valence and data type across polite F(3, 1148) = 3.01, $p = .031^{*}$ and appropriate F(3, 1148) = 2.71, $p = .045^{*}$ while considerate was trending F(3, 1148) = 2.62, $p = .051$. Thus, people's sensitivity to robot data collection is predicted by both the channel used *and* whether the data is used for positive or negative purposes.

Comment Addressee: Finally, as can be seen in Fig 6-3, comment addressee was also a very significant predictor of all labels: polite F(2, 1149) = 9.07, $p < .001^{**}$; considerate F(2, 1149) = 6.93, $p = .001^{**}$; appropriate F(2, 1149) = 6.80, $p = .002^{**}$, and for privacy F(2, 1149) = 18.75, $p < .001^{**}$. For example, comments that addressed both customers were rated as most polite, and excluding someone from being addressed by addressing one, or addressing one about the other resulted in *inconsiderate* and *inappropriate* ratings.

Meeting Variants: Meeting Valence was a significant predictor of considerate ratings F(1, 1150) = 8.06, $p = .005^{**}$, while Meeting Type was a significant predictor of appropriateness ratings F(2, 1149) = 4.40, $p = .013^{*}$. Meetings that are going badly lead participants to rate the robot as more inconsiderate, impolite, and inappropriate (Fig. 7-4). Participants also found robot comments made during the roommate meeting to be much more appropriate than those that occurred during a first date or job interview (Fig. 7-5).

Overall, social factors and data type played an important role in predicting participant ratings (even interacting with each other), with comment characteristics best predicting participant responses.

V. In-Person Study Results

One of the benefits of running an in-person study is that the research team can ask participants questions and view their faces (Fig. 8). This time, participant comments are included

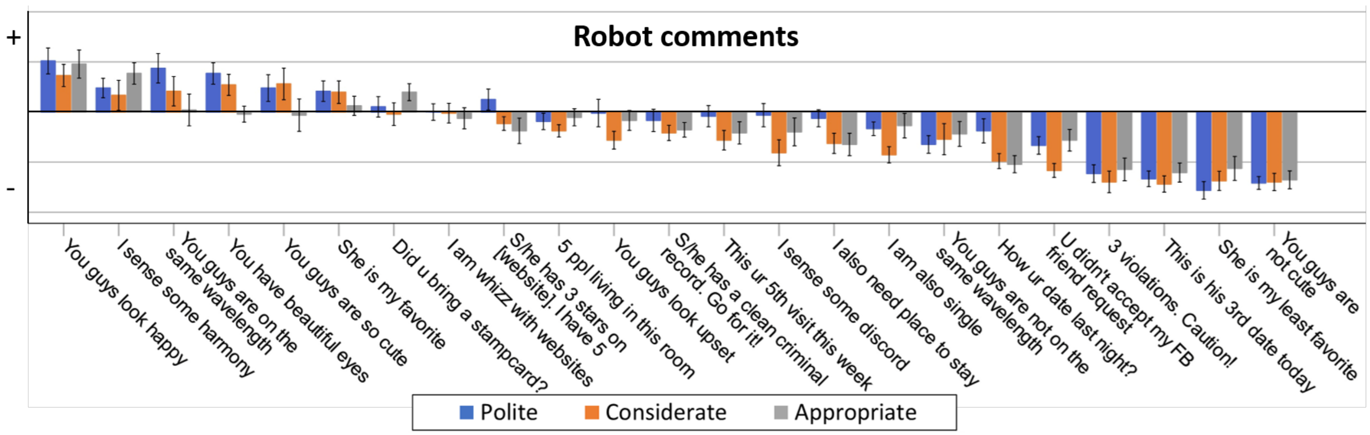

Fig. 5. Average rating of all robot statements in the online study.

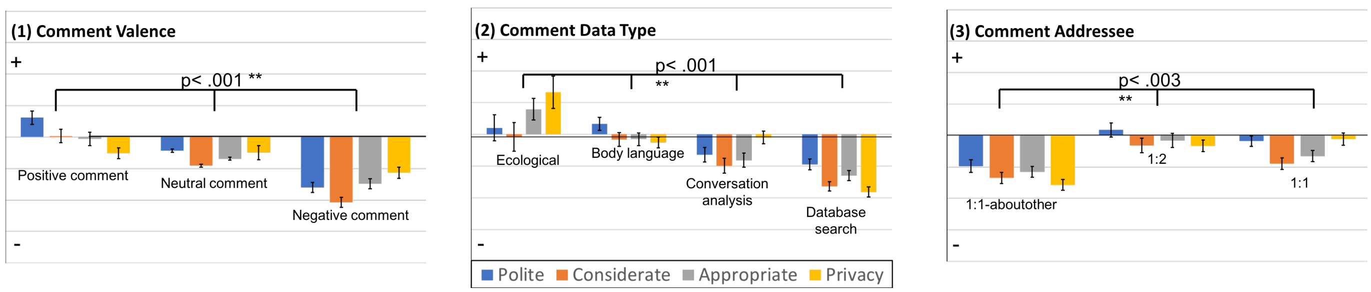

Fig. 6. ONLINE RESULTS: *Robot Comment Valence*, *Data Type*, and intended *Addressee* impact participant ratings. Positive comments are polite, while negative comments are highly inconsiderate. People are least sensitive to body language or cafe-related conversation, and are most reactive to comments in which the robot talks about one person to the other.

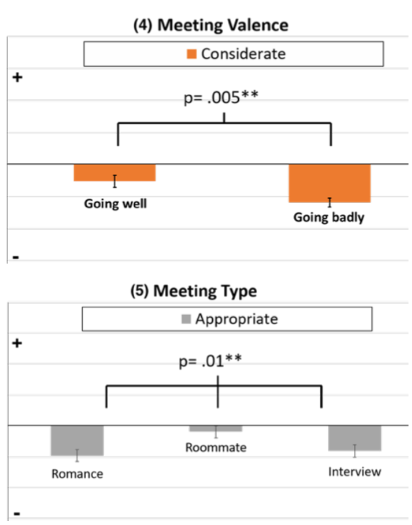

Fig. 7. ONLINE RESULTS (Meeting): (4) *Meeting Valence* predicts politeness ratings (4), i.e., it is impolite to comment during meetings that are going badly. Further, *Meeting Type* predicts appropriateness (5), i.e., comments are inappropriate during dates and interviews.

Fig. 8. Participants experienced and expressed emotions in the in-person study. Snapshots (a) and (b) demonstrate reactions to a negative robot comment, while (c) and (d) are to a positive robot comment.

alongside the Multi-factor ANOVA results and corresponding effect sizes (Table IV). Study conditions are presented in Table II, with the full set of robot statements in Fig. 4.

Comment Valence: Similar to the online study results, negative robot comments were rated negatively (Fig. 9-1). The effect occurred across all four labels, for polite F(2, 110) = 12.87, $p < .001^{**}$, considerate F(2, 110) = 10.82, $p < .001^{**}$, appropriate F(2, 110) = 7.38, $p < .001^{**}$, and privacy F()

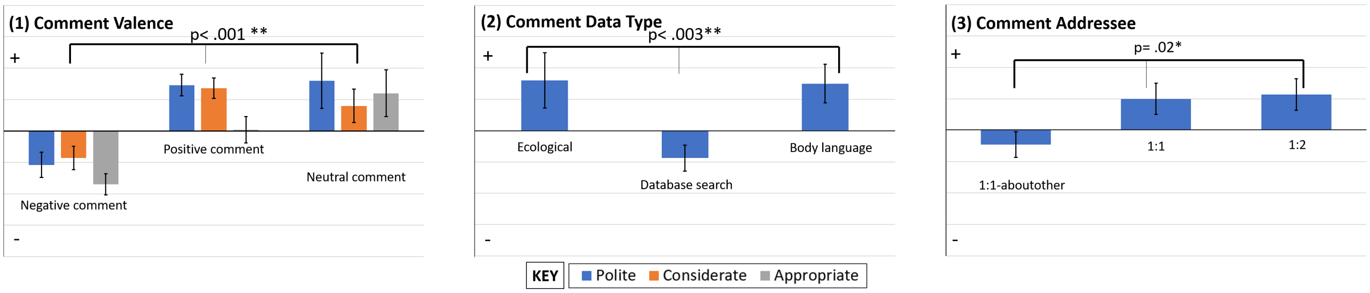

Fig. 9. IN-PERSON RESULTS:*Robot Comment Valence*, *Data Type*, and intended *Addressee* impact participant ratings. Positive comments are polite and considerate, while negative comments score low on all three labels. People are least sensitive to body language or cafe-related conversation, and are most disapproving of comments in which the robot talks about one person to the other.

= , p^{**}. For example, one participant responded to a negative robot comment with "not appropriate at all, the robot should be more professional". However, nice comments were perceived positively with one participant calling the robot "pretty chill, I think the robot is friendly and encouraging."

Comment Data Type: When analyzing the data type subgroup alone, there was a significant effect for polite F(2, 55) = 6.70, $p = .003^{**}$ but no effect for appropriateness and considerateness (Fig. 9-2). Again, we found positive responses to body-language related comments such as "You look happy together," with one participant writing, "Nice social robot." Database comments were again disliked, with one participant saying, "[the robot was] very kind for helping me but how does it know I have a clean criminal record?" But, again, responses to more. The Data Type comments also spurred conversation about social privacy, with one participant wondering whether "robots could eavesdrop," and another saying she did not mind the robot accessing her data, but "as a barista it was a little odd," which indicates people's potential openness to service-related data collection.

Comment Addressee: Addressee was a significant predictor for polite labels F(3, 109) = 2.73, $p = .020^{*}$, using the full set of in-person data. As depicted in Fig. 9-3, a robot addressing just one person was rated as impolite, while addressing one or both people was rated to be polite. One participant explicitly stated in the interview "when it addressed me it was better, even for the negative because it was at least to me; not like talking about me in front of me".

This time, *Meeting Valence* and *Meeting Type* had no significant results. Again, we find robot comment characteristics to be the best predictors of participant ratings.

Participants appreciated seeing multiple scripts, with one noting, "because we had 6 of them [I became] more comfortable with the emotions. And could concentrate on the actual interpretation." Interviews also included various projections of robot sensing, such as the robot reading the interviewers body language, scanning someones face and searching it in a database, or even detecting hesitations in a conversation. Participants suggested that the robot may have been trying to help break the ice in the first-date condition, but also that interrupting a job interview was always a bad idea. Exposing the participants to multiple conditions helps inspire these kinds of reflections, which expose additional variables for future exploration.

TABLE IV
IN-PERSON EFFECT SIZES

Label	Polite	Considerate	Appropriate
C.Valence	.236 (large)	.223 (large)	.109 (med.)
C.Data Type	.087 (med.)	.112 (medium)	.095 (med.)
C.Addressee	.070 (med.)	.033 (small)	.007 (none)

These interviews could also be used to seed future work. For example, one participant introduced the idea of personal history, suggesting that if the robot had a history with the people before that day, or was a coworker making coffee at the same workplace, that the range of comments the robot was making would become more appropriate. While another, quoted previously suggested that the data used should be appropriate to the robot's job, so perhaps future work can explore how personal history and job role further impact attitudes toward robot data use.

VI. DISCUSSION

The results reveal the social complexities of data privacy as it relates to social robots. Both the online and in-person studies underscored the relevance of situating robot data privacy research in a social context. In particular, while people have some opinions about what kind of data is permissible to use, most of the significant results came from the context in which that information is used:

- *Be nice:* The valence of the robot's comment strongly predicted social appropriateness. Overall, using data considerately was as important as the fact that the robot had collected a particular kind of data. This is not a typical consideration for data-security approaches, at least so far, but would be highly relevant to an interactive robot.
- *Reading body language is pretty innocuous:* Participants were most open to body language analysis, but were less likely to like conversation analysis, and hated being searched in databases. Future work should explore if this final finding would vary if the customer were offered

cafe-related perks or services, such as an automated punch card.

- *Try not to ignore people:* At least in the online study, people preferred the robot to address both people. Perhaps it is impolite to ignore someone when it is clear two people are there together. Or perhaps people react less to group comments than being singled out. This is a quality that could be explored in future work to clarify our findings.
- *Don't introduce traffic violations to dates or job interviews:* Overall, participants found robot comments least inappropriate during meetings between potential roommates. This indicates that romance and jobs are seen as higher stakes, so the robot should also know when to stay quiet.
- *Be especially nice if someone's already having a bad day:* Related to the above, if a robot can hear or see that someone's day is already going bad, it should assign more weight to using the data positively.

Several of the above points can be phrased as good manners, which reinforces the idea of socially-constructed privacy [26], previewed in Section 2.

The results also reveal interesting differences between what is appropriate, considerate, and/or polite, and when the terms converge: Positive statements about any topic and body language analysis were rated as *polite*. Cafe-related conversation was deemed most *appropriate*. And while few statements were rated *considerate*, there were many inconsiderate ratings for negative comments, database search, and conversation analysis.

Overall, we found this format – using generative scripts to explore many different privacy features – to be effective and thought-provoking. While these results do not guarantee we have captured the nuances of how real customers might respond to a service robot's attempts at casual or data-based conversation, they do provide a starting point, and allow us to explore privacy concepts and violations that would be unethical to explore otherwise.

VII. Conclusion

This study investigated social attitudes towards a robot using their data conversationally. We conducted two studies: first, an online study (N=4608) where the participants rated a video of an interaction of a robot barista with two people, and second, an in-person (N=20) study where the participant was one of the actors in the scene.

Both studies found that using the data positively was as important as the type of data that was used. Unlike previous approaches to data privacy, this suggests social factors play a *primary* role in human acceptance of data collection by a robot. The online study offered us a chance to look at how people would perceive a robot interacting with other customers, while the in-person study showed how people would like the robot to interact with them.

Using a theater-inspired method allowed participants to experience a variety of scenes without putting their personal privacy at risk. In fact, in-person participants reported feeling real emotions during the scenes. Most online study results were replicated in the in-person study, and the dual format allowed us to explore many variables (online), and also ask questions of people (in-person).

In future work, we would like to explore how other types of robot comments and situated factors influence perceptions of a robot barista. It may be fruitful to conduct a naturalistic study in a real cafe, a setting with more ecological validity. It appears that service-related conversation is the lowest risk, however, politeness ratings may benefit from greater intimacy, as long as it is positive or based on non-verbal signals. An ecological study would also establish how well these simulated findings map to the real world. It would be interesting to assess if how these results might vary by culture.

Context is extremely important to robot data use. How the meeting between the customers is going affects the perception of service robot social appropriateness. People may care more about their social reputation on a job interview than one a data, and robots should definitely be cautious about talking to one customer about another. Across all of our data, however, one finding is clear: robots of the world – whatever your data perception capabilities – frame your knowledge positively!

VIII. Acknowledgements

We would like to thank our actors Jessica Andrade and Melissa Santala, and lab members Alexandra Bacula, Abrar Fallatah, and Jeremy Urann for assisting with the studies. We would also like to thank Imagine Coffee in Corvallis where we filmed the videos for the online study.

References

[1] "Robot 'hired' by scottish supermarket then 'sacked' after a week because he couldn't understand customers," *Daily Record*, 2018. [Online]. Available: https://www.dailyrecord.co.uk/news/science-technology/robot-hired-supermarket-sacked-after-11893730
[2] D. Solove, *Understanding privacy*. Harvard University Press, 2008.
[3] E. T. Hall, *The Hidden Dimension*. Garden City, N.Y., Doubleday, 1966.
[4] H. Nissenbaum, "Privacy as contextual integrity," *Wash. L. Rev.*, vol. 79, p. 119, 2004.
[5] E. A.Allen, E. N. Zalta, "Privacy and medicine," *The Stanford Encyclopedia of Philosophy*, Spring 2011.
[6] J.-W. Byun, E. Bertino, and N. Li, "Purpose based access control of complex data for privacy protection," in *Proceedings of the tenth ACM symposium on Access control models and technologies*. ACM, 2005, pp. 102–110.
[7] K. A. Bamberger and D. K. Mulligan, *Privacy on the ground: driving corporate behavior in the United States and Europe*. MIT Press, 2015.
[8] R. L. Krutz and R. D. Vines, *Cloud security: A comprehensive guide to secure cloud computing*. Wiley Publishing, 2010.
[9] S. Preibusch, "Guide to measuring privacy concern: Review of survey and observational instruments," *International Journal of Human-Computer Studies*, vol. 71, no. 12, pp. 1133–1143, 2013.
[10] S. Yang, B. K.-J. Mok, D. Sirkin, H. P. Ive, R. Maheshwari, K. Fischer, and W. Ju, "Experiences developing socially acceptable interactions for a robotic trash barrel," in *Robot and Human Interactive Communication (RO-MAN), 2015 24th IEEE International Symposium on*. IEEE, 2015, pp. 277–284.
[11] C. Nass, J. Steuer, and E. R. Tauber, "Computers are social actors," *Conference companion on Human factors in computing systems - CHI '94*, p. 204, 1994. [Online]. Available: http://portal.acm.org/citation.cfm?doid=259963.260288

[12] C. Breazeal, "Social interactions in hri: The robot view," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 34, no. 2, pp. 181–186, 2004.

[13] M. R. Calo, "12 robots and privacy," *Robot ethics: The ethical and social implications of robotics*, p. 187, 2011.

[14] M. Rueben, C. M. Grimm, F. J. Bernieri, and W. D. Smart, "A taxonomy of privacy constructs for privacy-sensitive robotics," 2017.

[15] M. Rueben, F. J. Bernieri, C. M. Grimm, and W. D. Smart, "Framing effects on privacy concerns about a home telepresence robot," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '17. New York, NY, USA: ACM, 2017, pp. 435–444. [Online]. Available: http://doi.acm.org/10.1145/2909824.3020218

[16] M. K. Lee, K. P. Tang, J. Forlizzi, and S. Kiesler, "Understanding users' perception of privacy in human-robot interaction," in *Proceedings of the 6th international conference on Human-robot interaction*. ACM, 2011, pp. 181–182.

[17] L. Barkhuus, "The mismeasurement of privacy: using contextual integrity to reconsider privacy in hci," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2012, pp. 367–376.

[18] T. Denning, C. Matuszek, K. Koscher, J. R. Smith, and T. Kohno, "A spotlight on security and privacy risks with future household robots: attacks and lessons," in *Proceedings of the 11th international conference on Ubiquitous computing*. ACM, 2009, pp. 105–114.

[19] Y. Moon, "Intimate exchanges: Using computers to elicit self-disclosure from consumers," *Journal of consumer research*, vol. 26, no. 4, pp. 323–339, 2000.

[20] D. B. Shank, "Are computers good or bad for business? how mediated customer–computer interaction alters emotions, impressions, and patronage toward organizations," *Computers in Human Behavior*, vol. 29, no. 3, pp. 715–725, 2013.

[21] A. Reben and J. Paradiso, "A mobile interactive robot for gathering structured social video," in *Proceedings of the 19th ACM international conference on Multimedia*. ACM, 2011, pp. 917–920.

[22] S. Kiesler, A. Powers, S. R. Fussell, and C. Torrey, "Anthropomorphic interactions with a robot and robot–like agent," *Social Cognition*, vol. 26, no. 2, pp. 169–181, 2008.

[23] N. Martelaro, V. C. Nneji, W. Ju, and P. Hinds, "Tell me more: Designing hri to encourage more trust, disclosure, and companionship," in *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 2016, pp. 181–188.

[24] P. Lin, K. Abney, and G. A. Bekey, *Robot ethics: the ethical and social implications of robotics*. The MIT Press, 2014.

[25] P. M. Asaro, "Robots and responsibility from a legal perspective," *Proceedings of the IEEE*, pp. 20–24, 2007.

[26] I. Altman, "Privacy Regulation: Culturally Universal or Culturally Specific?" *Journal of Social Issues*, vol. 33, no. 3, pp. 66–84, 1977.

[27] J. K. Burgoon, "Privacy and communication," *Annals of the International Communication Association*, vol. 6, no. 1, pp. 206–249, 1982.

[28] I. Altman, *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding.* ERIC, 1975.

[29] H. Nissenbaum, *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2009.

[30] M. Haahr, "RANDOM.ORG: true random number service," https://www.random.org, 1998–2018, accessed: 2018-06-01.