

A Touching Connection: How Observing Robotic Touch Can Affect Human Trust in a Robot

Theresa Law · Bertram F. Malle · Matthias Scheutz

Received: date / Accepted: date

Abstract As robots begin to occupy our social spaces, touch will increasingly become part of human-robot interactions. This paper examines the impact of observing a robot touch a human on trust in that robot. In three online studies, observers watched short videos of human-robot interactions and provided a series of judgments about the robot, which either did or did not touch the human on the shoulder. Trust was measured using a recently introduced multi-dimensional instrument, which assesses people’s trust in a robot as being capable, reliable, sincere, and/or ethical. The first study showed that observed robot touch increased overall trust in the robot, especially for the sincere and ethical trust aspects, and led people to perceive the robot as more comforting, but also more inappropriate. A second study replicated the general pattern, even with a handshake preceding the touch; but in the context of the handshake the touch was seen as more inappropriate. A third study examined the joint impact of a handshake, touch, and information about the robot’s designed function. In the context of such information, observed touch was seen as even more inappropriate, which in turn decreased trust.

Keywords Robot touch, human-robot trust, social connection, human-robot interaction

T. Law · M. Scheutz
HRI Lab
Tufts University
Medford, MA 02155
E-mail: {theresa.law, matthias.scheutz}@tufts.edu

B. F. Malle
Brown University
Providence, RI 02912
E-mail: bfmalle@brown.edu

1 Introduction

As robots expand into the human social world, previous physical barriers such as safety cages have disappeared. As a result, people are able to touch robots and robots may touch people—accidentally, or even intentionally. Because human-human touch can have nuanced, even charged meanings, it is important to understand how people react to robots occupying roles in which they might initiate touch. Tactile HRI has already begun to be established with therapy robots such as *Paro*, which thus far have yielded promising results for elevating people’s mood and alleviating stress [42] (for a recent review, see [24]). However, less research is available on how robot-initiated touch may affect people. For the healthcare domain in particular, robot-initiated touch will be common and necessary—such as when robots take vitals, perform physical therapy, or help people out of bed. Beyond the tactile feeling of being touched, people’s psychological reactions to a robot’s touch may vary by the person’s gender, the robot’s perceived gender, the person’s social status, the setting of the touch, and many other complicating factors that have been examined in the human-human touch literature [18].

Before even interacting with a robot that could initiate touch, a person often has the opportunity to observe how the robot interacts with others. This situation happens regularly in everyday life: We observe a person interacting with someone else, and that interaction influences how we perceive the person before we meet them ourselves. While impressions of an agent can carry over from observation to interaction, they are not guaranteed to match. For example, Torrey et al. [51] asked people to watch a video of a robot giving instructions to a human, and people rated a robot that

gave indirect hedged commands as more likeable than a robot that gave direct commands. However, when researchers attempted to replicate these results with both an observation-based study and an in-person study [50], they found that, while people preferred the indirect robot in the observational paradigm, they preferred the direct robot in the in-person interaction. In human neurophysiology, too, despite previous suggestions that merely observing human action can activate the somatosensory cortex (hence produce an analogous experience of touch) [6], recent studies have questioned this claim [10]. Thus, even though observation and experience of touch share important similarities, they are distinct; especially in the case of robot touch, observational paradigms and in-person interaction paradigms present separate but complementary perspectives.

Among the core responses to social robots, trust has received increasing attention of late. Previously robots had been studied primarily as being more or less reliable machines [21], but their increasing presence in social contexts and their social-interactive capacities make them inevitable targets of the more relational aspects of trust [29, 33]. Given the potentially relational meaning people impose on robot-initiated touch, the aim of this paper is to explore how observing robot-initiated touch affects how much a human observer trusts the robot. We conducted three online studies that relied on an observation-style paradigm. These studies begin to document how perceived robot behavior affects a person's impressions of the robot prior to ever interacting with it. We do not claim that these findings straightforwardly generalize to human perceptions during or after interacting with a robot. But the inferences people draw from mere observation are already so complex as to warrant systematic investigation.

We begin with a brief overview of touch in the human-human interaction literature, then review touch in the human-robot interaction literature, and finally trust in the human-robot interaction literature. We then outline the hypotheses and report on Experiment I. The findings of Experiment I informed the methodology of Experiment II, which, in turn, guided the methodology of Experiment III. We close with general insights gained from these experiments and point to limitations and future directions.

2 Background and Related Work

2.1 Touch

2.1.1 Human-Human Interaction

In the human-human interaction literature, touch is widely seen as having positive, pro-social effects. A small touch to the shoulder or hand has been found to increase restaurant tipping [13], shopper compliance [23], and having favorable impressions of people [15, 16]. This phenomenon has been referred to as the *Midas Touch* [13]. However, the meaning of touch is delicate, and complicated by a number of factors. Which parts of the body are permitted to be touched depends on the toucher's and recipient's gender, their relationship, and religious affiliations [27]. Different settings can result in different meanings of touch [26], and acts of touching have both explicit and implicit meanings [5]. Touch is also a type of interaction that is particularly sensitive to individual differences. While some people react very positively to touch, others naturally avoid touch and react negatively [1].

2.1.2 Human-Robot Interaction

Touch in HRI is an under-explored field. Much of the early work focused on observation studies exploring how different populations naturally touched robots [46, 45, 44]. Subsequent experimental studies documented beneficial results when people initiate touch with a robot. In fact, participants preferred to touch a robot over being touched by it [20], and when they do touch a robot, people convey emotions to the robot in a similar manner as they do to humans [2]. In a study in which participants watched a movie with a Pleo robot, participants perceived their friendship with the robot and their own emotional stability as higher when they touched the robot during the movie [38]. Another study, however, found no difference in participant's perceived stress levels as a function of touching the robot [7].

Studies in which robots initiate touch show mixed results. In one study [11], participants acted out a healthcare scenario in which a robot nurse performed a comforting or instrumental touch on the participant's arm, and either did or did not give a verbal warning that the touch was about to happen. Participants most preferred the instrumental touch with no warning condition and generally found the instrumental touch to be more enjoyable and necessary than the comforting touch. Research has also shown that active robot touch,

such as when a robot strokes a person’s hand while they complete a monotonous task, led people to continue working on the task longer [47,34]. By contrast, other researchers found no impact of a robot holding participants’ hands while watching a scary movie [55]. A hug from a large robot teddy bear encouraged pro-social behavior in people, such as donating more money to charity [48]. However, another study did not find that robot-initiated touch increased pro-social behaviors, but it did lower physiological stress and strengthened the perceived human-robot bond, whether or not a social bond had previously been established [56]. One study found that touch increased positive evaluations of the robot [3], but these results were complicated by differences due to participant gender. A set of unpublished studies found numerous moderators of perceived acceptance of a robot’s touch, including human attitudes, robot appearance, type of touch, and more [22]. In many previous studies [11,20,34,47,48], the robot first asks for permission or offers some other verbal cue before initiating touch. In everyday life, touch is more likely to occur spontaneously without a verbal warning, so in our studies we focused on this form of touch.

2.2 Trust

Trust involves an agent being willing to be vulnerable to the actions of another in a situation marked by uncertainty [30]. Any human-robot interaction, therefore, raises questions of trust. Human trust in robots can be divided into two categories—performance-based trust, in which the human trusts that the robot is capable of completing a task, and relational or moral trust, in which the human trusts that the robot has knowledge of social norms [29,33]. Both of these types of trust may be affected by a robot touching a human: People need to trust that the robot is physically capable of not harming them and that it understands when it is socially appropriate to initiate touch.

At least in humans, touch can indicate trust between two agents. Findings show that photographs of two people who touch each other convey more receptivity and trust to a viewer than corresponding photographs without touch [9].

Few studies have looked at the effects of touch on trust in HRI. In one study [17], researchers attempted to replicate the Midas Touch effect seen in human-human interaction [13] by having participants play a behavioral economics game with a robot while wearing an EEG net. In the game, the robot would make unfair monetary offers to

the participant and either stroke the participant’s hand as it made the offer or not. The researchers found that when the robot stroked the person’s hand, people’s feelings of unfairness were inhibited at the level of EEG activity. However, behaviorally, there was no difference between the touch and no touch conditions: People accepted or rejected the robots offers at equal rates. Other researchers found that participants who themselves touched an android robot’s arm during an economic bargaining game showed behavioral decisions that reflected more trust in the robot [14]. Likewise, people who watched a horror film clip while holding a robot’s warm hand increased feelings of friendship and trust toward the robot [36].

In this article, we operationalize trust using the Multidimensional Measure of Trust (MDMT), a recently developed brief trust questionnaire in which participants rate how much they feel that a robot has certain properties of trustworthiness [53,52]. Critically, the MDMT separately measures performance aspects of trust (subscales of *Capable* and *Reliable*) as well as moral aspects of trust (subscales of *Sincere* and *Ethical*) [54,33]. Ascribing moral trustworthiness to an agent may be particularly relevant when examining an observer’s impression of the agent’s touch. In our studies, we examine both overall trust levels (average across all subscales) as well as the two components of performance trust (*Capable* & *Reliable*) and moral trust (*Sincere* & *Ethical*).

3 The Current Experiments

In Experiment I, we examined how participants’ trust in a robot is affected by observing the robot touch another person on the shoulder during a brief video-recorded interaction. However, previous research suggested that a robot’s attitude and the human actor’s gender can influence participants’ impressions of a human-robot interaction [3,22]. To explore these potential moderators and increase generalizability we therefore varied the robot’s attitude and the human’s gender. We also varied the robot’s “gender” as this variable has been shown to exert a significant influence on people’s perceptions [28,37,43] and might interact with the human’s gender. To indicate robot gender, we varied the robot’s voice. Synthetic voices have been shown to powerfully influence gender perceptions, with males and females trusting male and female synthetic voices differently [35]. Gendered robot voices have also been shown to influence human behavior in human-robot interaction [40,43,12], as well as perceptions of robot trustwor-

thiness [28]. We hypothesized that observed touch will result in perceptions of the robot being comforting (the conventional interpretation of such a gesture) and increase trust, though we expected moderation of this effect by different gender combinations of robot and human (e.g., participants may see a male robot touching a female human actor to be less appropriate).

In Experiment II, we explored whether the same observed shoulder touch would continue to affect observers' trust in the robot when the touch is displayed in the context of a previously established social connection between the robot and the human actor. To examine this question, we presented an image of the human and robot shaking hands before the video began and varied whether the video displayed robot touch or not. Shaking a robot's hand has been shown to increase overall favorable ratings of a robot [4] and, among humans, a handshake tends to have positive effects on trust [9, 31]. We therefore expected that the handshake would set a positive baseline for trust in the robot and examined whether the shoulder touch would make an additional contribution to trust or even mitigate some negative perceptions we had seen in Experiment I.

In Experiment III, we examined the impact of the shoulder touch relative to another context factor: whether the robot was designed to be customer-focused or performance-focused. Participants either did or did not see the handshake photo, did or did not see the shoulder touch in the video, and were told either nothing about the robot's design or that it was designed to be customer-focused or performance-focused. We thus wanted to test again the patterns of findings in Experiments I and II and examine whether additional information about the robot's function moderates the interpretation of touch and its impact on trust. An exploratory question was whether a customer-focused role has differential impact on moral trust and a performance-focused role has differential impact on performance trust.

All study procedures were approved by our institution's IRB (Protocol #1609030).

4 Experiment I

4.1 Methods

4.1.1 Participants

We planned to recruit 600 participants, 25 in each cell of the 2 (Touch/No Touch) $\times$ 3 (Positive/Neutral/Negative Attitude) $\times$ 2 (Male/Female Robot)

$\times$ 2 (Male/Female Actor) design. After eliminating empty or incomplete records, a total of 587 people remained in the sample (with 23 to 25 participants per cell). We recruited participants through Amazon Mechanical Turk (AMT). Self-reported gender was 260 female, 325 male, 2 no answer. Their ages ranged from 19 to 73 ($M = 35.6$, $SD = 10.46$). The ethnic composition was: 73.5% White or Caucasian, 10.2% African American, 7.04% Hispanic, 6.1% Asian, 1.4% American Indian or Alaska Native, 0.3% Native Hawaiian or Pacific Islander, 1.3% Other or Prefer Not to Answer.

4.1.2 Materials

All experiments reported here presented participants a video of a brief human-robot interaction in which, by random assignment, the robot either touched a human on the shoulder or not. We chose a video over photographs to increase the natural dynamics of interaction while maintaining full stimulus control. As the touch manipulation, we selected a shoulder touch for two reasons. Previous research found that people perceive a person's shoulder to be an acceptable body part for a robot to touch [22], and we assumed that a shoulder touch allows for a relational interpretation that can affect perceptions of trustworthiness.

The video opened with the human standing at a computer screen and the robot (a Willow Garage PR2) standing on the left side of the human, angled so that it was facing both the human and the screen (Fig. 1). As the human entered information on the computer, the robot uttered phrases like "alright" and "okay". After a few seconds, the computer screen went dark, and the human turned to look at the robot. The robot responded with an utterance based on the *Attitude* condition. While speaking, it performed a behavior depending on the *Touch* condition: It either touched the human on the back of the shoulder with its right hand while gesturing to the keyboard with its left hand, or it solely gestured to the keyboard with its left hand. If the robot touched the human's shoulder, its hand stayed there for about 1.5 seconds before returning to its side. We chose the durations of each phase of the video so as to make the sequence seem natural and realistic, using the judgments of about a dozen members of our research groups. The robot then told the human how to fix the error. The human did what the robot said and continued to enter information into the computer for a few seconds before the video faded out. Each video lasted approximately 20 seconds—long enough to set a realistic context for the key actions but short enough to maintain participants' atten-

tion and make the touch sufficiently noticeable. The manipulated and measured between-subject factors were as follows:

- **Touch:** after the computer crashed, the robot either did or did not touch the back of the human’s shoulder.
- **Attitude:** when the computer crashed, the robot responded in one of the three ways:
 - *Positive:* “You did not do anything wrong. Leave it to me. Press escape twice.”
 - *Neutral:* “What could have gone wrong. Let me see. Press escape twice.”
 - *Negative:* “What have you done wrong. Listen to me. Press escape twice.”
- **Robot Gender:** the robot spoke with either a male voice or a female voice. The voices were created with the MaryTTS online platform using the “cmu-slt-hsmm en US [female/male] hmm” options.
- **Human Actor Gender:** the human at the computer was either male or female.




Fig. 1 (a) Male Human, No Touch condition. (b) Female Human, Touch condition.

4.1.3 Measures

Trust. To assess participants’ trust in the robot, we used the Multidimensional Measure of Trust (MDMT) [53,33], which asks participants to rate, on 0-7 point scales, how well 16 descriptors apply to the robot. These words fall into four groups, based on validation studies [54], and form four subscales: Reliable, Capable, Sincere, and Ethical (see Figure 2 for individual items in each subscale). The initial version [53] showed reliability (Cronbach’s alpha) values of .72 to .88; a revised version, used here, found reliabilities of 0.79 to 0.92 [54]. This study also showed that at least the combined pairs of Reliable & Capable and Ethical & Sincere are responsive to relevant manipulations of information about robot behavior. To address the limited validation so far, the authors have made the measure publicly available on their website and encourage other researchers to use it and help compile the findings into a multi-researcher validation project.

Behavior impressions. A brief look into various body language internet resources suggests that a shoulder touch is typically nonsexual but can still be interpreted in various ways [8]. To explore people’s interpretations of a robot touching a human, we presented six adjectives: comforting, measured, surprising, odd, inappropriate, and creepy. Several of these adjectives reflected spontaneous descriptions that participants gave to a free-response item in a pilot study; the word *measured* was added to make the set contain three positive and three negative items. Participants used 0-7 point rating scales to express how well each of these adjectives described the robot.

Sincere	Ethical	Capable	Reliable
Sincere	Ethical	Capable	Reliable
Genuine	Respectable	Skilled	Predictable
Candid	Principled	Competent	Count on
Authentic	Has integrity	Meticulous	Consistent

Fig. 2 Multi-dimensional trust as assessed by the MDMT, with its four dimensions (subscales) of trust and the four words used to assess each dimension.

4.1.4 Procedure

Participants were recruited from Amazon Mechanical Turk for an online study. After providing informed consent, participants were told that they would be watching a few short videos of a variety of human-robot interactions in order to study people’s impressions of robots. All participants first watched a warm-up video, in which a person controls a robotic wheelchair. They then answered a series of questions about their impressions of the robotic wheelchair. After the wheelchair questions, participants watched one of the 24 videos, corresponding to their randomly assigned condition. They were not allowed to skip, pause, or rewatch the video. After the video ended, they completed the MDMT trust scales, prompted with the phrase “Please consider the following descriptors and indicate how well they apply to the robot.” Thereafter, they answered questions about the robot’s mental capacities (not reported here). Next they were asked to “please tell us about a few more impressions you had of the robot: How did the robot’s behavior seem to you?”, which was followed by the six behavior impression questions. Finally, participants answered a manipulation check question and provided demographic information. The manipulation check asked if the participant saw the robot touch the person; response options were

“yes”, “no”, or “maybe”. Participants who answered this question incorrectly (“no” in the touch condition or “yes” in the no-touch condition) were excluded; participants who said “maybe” were included. To minimize demand characteristics (inviting people to say “yes”) the question about touch was preceded by two seemingly similar questions: whether participants saw the robot move forward and whether they saw it make a gesture. Responses to these questions were not analyzed.

4.2 Results

Forty-nine participants (8.2%) failed the manipulation check, leaving 538 for analysis. We conducted between-subjects ANOVAs on the 2 (Touch/No Touch) $\times$ 3 (Positive/ Neutral/ Negative Attitude) $\times$ 2 (Male/Female Robot) $\times$ 2 (Male/Female Actor) design for each dependent variable.

4.2.1 Trust

Analyzing overall trust, we found a significant main effect of *Touch* $F(1, 514) = 6.9, p = .009$, Cohen’s $d = 0.23$. Participants who saw the robot touch the human ($M = 4.63, SD = 1.43, n = 278$) trusted the robot more than those in the no-touch condition ($M = 4.30, SD = 1.47, n = 260$). No other effects were significant at $p < .05$.

For the **Sincere & Ethical** trust component specifically, we again found the main effect of *Touch*, $F(1, 514) = 9.7, p = .002, d = 0.27$. Participants who saw the robot touch the human ($M = 4.40, SD = 1.62, n = 278$) trusted the robot more than those in the no-touch condition ($M = 3.96, SD = 1.64, n = 260$). No other effects in the tested design were significant.

For the **Capable & Reliable** trust component, the *Touch* effect was in the expected direction but weaker ($d = 0.16$) and did not reach traditional significance, $F(1, 514) = 3.2, p = .075$. In addition, a three-way interaction emerged between *Robot Gender*, *Actor Gender*, and *Robot Attitude* $F(2, 514) = 4.0, p = .019, \eta_p^2 = .015$. However, Tukey HSD post-hoc tests revealed no specific significant comparisons.

4.2.2 Behavior Impressions

Correlations among the behavior descriptors suggested a considerable affinity among Inappropriate, Odd, and Creepy ($r_s > .60$), with Surprising somewhat related to them ($r_s = .19$ to $.41$). Comforting and Measured each stood on their own.

Confirming these observations, a Principal Component Analysis (PCA) with orthogonal rotation yielded two components. The first combined Inappropriate, Odd, and Creepy, the second was primarily constituted by Comforting, while Surprising loaded on both components and Measured on neither. We, therefore, formed a composite of Inappropriate, Odd, and Creepy and analyzed the other three descriptors on their own.

For **Comforting**, a main effect of *Touch* emerged, $F(1, 514) = 38.3, p < .001, d = 0.54$. Participants who saw the robot touch the human ($M = 4.10, SD = 2.04, n = 278$) rated the robot’s behavior as more comforting than those in the no-touch condition ($M = 2.95, SD = 2.21, n = 260$). There was also a main effect of *Attitude*, $F(2, 514) = 5.9, p = .003, d = .31$. A Tukey HSD posthoc test revealed that participants rated the robot’s behavior as more comforting when the robot displayed a positive attitude ($M = 3.83, SD = 2.10, n = 185$) or a neutral attitude ($M = 3.69, SD = 2.06, n = 181$) as compared to a negative attitude ($M = 3.09, SD = 2.37, n = 172$).

For the **Inappropriateness** composite, there was a main effect of *Touch* $F(1, 514) = 3.75, p = .053, d = 0.17$. Participants who saw the robot touch the human rated the robot’s behavior as more inappropriate ($M = 2.82, SD = 1.96, n = 278$) than those in the no-touch condition ($M = 2.52, SD = 1.80, n = 260$). There was also a main effect of *Attitude* $F(2, 514) = 7.83, p < .001, d = .35$. A Tukey HSD posthoc test revealed that participants rated the robot’s behavior as significantly more inappropriate when the robot displayed a negative attitude ($M = 3.13, SD = 1.97, n = 172$) than when it displayed a neutral ($M = 2.49, SD = 1.84, n = 181$) or positive attitude ($M = 2.42, SD = 1.79, n = 185$).

For **Measured**, there was no effect of *Touch*, $d = 0.01$. A three-way interaction emerged among *Touch*, *Robot gender*, and *Attitude*, $F(2, 514) = 4.71, p = .009, \eta_p^2 = .017$, but Tukey HSD posthoc tests revealed no robust differences.

Finally, for **Surprising**, there was a significant main effect of *Touch* $F(1, 514) = 24.6, p < .001, d = 0.43$, with participants who saw the robot touch the human finding the robot’s behavior more surprising ($M = 3.96, SD = 2.10, n = 278$) than those in the no-touch condition ($M = 3.04, SD = 2.10, n = 260$).

4.2.3 Multivariate Analyses

The experimental manipulation of touch had effects on trust and on several behavior impressions, both positive ones (Comforting), negative

ones (Inappropriate), and neutral ones (Surprising). We explored how these effects related to each other. A plausible pattern is that the robot’s touch affects people’s interpretations of the robot’s behavior (as more comforting but also more surprising and inappropriate), which, in turn, affect the felt trust in the robot. We tested this model using the JASP [25] mediation module. The initial effect of *Touch* on overall trust was positive and significant (unstandardized $b = 0.32, z = 2.6, p = .009$), but seeing the robot as being comforting was a significant mediator ($b = 0.46, z = -5.86, p < .001$), eliminating the direct effect ($b = -0.13, z = -1.29, p = .199$). This mediation pattern was stronger for Sincere & Ethical than for Capable & Reliable. Thus, while *Touch* has an effect on interpretations of being comforting and being perceived as comforting predicts trust, the touch action makes no independent predictive contribution to rated trust.

The other behavior impressions did not function as significant mediators between touch and trust. However, the set of behavior impression variables (without the *Touch* manipulation) strongly predicted trust in a multiple regression analysis, $R = .70, F(4, 533) = 127.4, p < .001$. Over and above Comforting ($r_{sp} = .38$), we find significant predictive contributions from Measured ($r_{sp} = .31$), the Inappropriate composite ($r_{sp} = -.24$), and Surprising ($r_{sp} = .11$), all $ts > 3.0, ps < .001$.

4.3 Discussion

Observing a robot touch a human on the shoulder affected people’s impressions of the robot’s behavior and, in turn, their trust in the robot. The MDMT trust component that focuses specifically on social-relational and moral qualities (Sincere & Ethical) was more sensitive to this influence than the performance-focused trust component (Capable & Reliable). The impression that the robot’s touch elicited came with an interesting tension between being perceived as comforting (the more immediate meaning that such a gesture normally implies) but also as somewhat odd and inappropriate.

We did not find robust trust effects due to robot attitude or any interactions between robot gender and human gender, as others have found [3]. The expected effects of the robot’s attitude on behavior impressions (comforting and the inappropriateness composite) are unsurprising but lend internal validity to the measures and manipulation.

In light of these results, our next experiment explored how context affects people’s interpretation of the robot’s touch, and in turn their trust

in the robot. The tension between impressions of the robot being comforting and inappropriate may stem from the fact that there is no background information on the robot’s and the human’s role, or how the robot and human are socially connected. In Experiment II, participants were, therefore, shown a picture of the human and robot shaking hands before the main interaction video, which then did or did not contain the touch. This picture was intended to establish a social connection between the robot and the human, thus setting a positive baseline for likability [4] and trust [9,31]. Having inferred a social connection from the handshake and having developed a reasonably positive impression from it, people may resolve the mixed impression of the shoulder touch in favor of being comforting and against being inappropriate.

Because of a clerical error we ran the touch/no-touch conditions only with the handshake displayed in the beginning of the video, omitting the no-handshake control condition. However, we conducted a cross-experiment comparison between the corresponding conditions of Experiment II with those of Experiment I that had no handshake photo, and we implemented a fully crossed design in Experiment III.

5 Experiment II

5.1 Methods

5.1.1 Participants

A total of 60 people participated in this study through Amazon Mechanical Turk (female: 15, other: 1). Their ages ranged from 19 to 70 ($M = 37.1, SD = 11.05$). The ethnic composition was as follows: 62.5% White or Caucasian, 18.8% Black or African American, 9.4% Asian, 4.7% Hispanic, 1.6% American Indian or Alaska Native, 1.6% Native Hawaiian or Pacific Islander, 1.6% Other.

5.1.2 Procedure and Materials

The procedure was the same as in Experiment I. However, in the introduction to the main video, participants saw an image of the human and the robot shaking hands (Fig. 3). Then they saw the main video of the brief interaction showing a female human actor and a male robot with neutral attitude. The neutral attitude version was chosen because we wanted to work at the baseline of impressions. Further, in the absence of robust effects of human and robot gender in Experiment I, any combination of the gender variables would

have qualified, but we felt that the female human with male robot ensured enough room for varied interpretations of the touch. The measures were the same as in Experiment I: the MDMT-based trust ratings and the six behavior impression ratings, followed by the manipulation check and demographic questions.

5.2 Results

Nine participants (15%) failed the manipulation check question, leaving 51 participants for analysis.

5.2.1 Trust

We first examined the effect of *Touch* vs. *No Touch* on overall trust and also separately on Sincere & Ethical and Capable & Reliable trust. Though all the means were in the same direction as in Experiment I (see Table 1), none of them reached traditional significance levels. For overall trust, $F(1, 49) = 1.51$, $p = .22$, $d = 0.34$. For Sincere & Ethical, $F(1, 49) = 1.75$, $p = .19$, $d = 0.37$. For Capable & Reliable, $F(1, 49) = 0.36$, $p = .55$, $d = 0.17$. Note that the effect sizes in Experiment II were in fact larger than those in Experiment I (e.g., $d = 0.37$ vs. 0.27 for Sincere & Ethical).



Fig. 3 Image seen before the video in the *Handshake* conditions.

5.2.2 Behavior Impressions

As in Experiment I, the descriptors Odd, Inappropriate, and Creepy were substantially correlated and formed a clear principal component, whereas Comforting formed its own component, Surprising loaded on both, and Measured loaded on neither. Thus, we analyzed an **Inappropriate** composite as well as the remaining descriptors separately. Compared to the no-touch control, the

Table 1 Means (& SDs) for Experiment I (No Handshake) and Experiment II (Handshake) (* denotes values in touch conditions that were significantly higher than their no touch counterpart)

		Exp I: No Handshake	Exp II: Handshake
All Trust	Touch	4.63 (1.43)*	4.85 (1.24)
	No	4.30 (1.47)	4.41 (1.32)
Sincere & Ethical	Touch	4.40 (1.62)*	4.65 (1.31)
	No	3.96 (1.64)	4.07 (1.76)
Capable & Reliable	Touch	4.85 (1.47)	4.97 (1.28)
	No	4.85 (1.38)	4.76 (1.21)
Comforting	Touch	4.10 (2.04)*	4.77 (2.03)*
	No	2.95 (2.07)	2.92 (1.91)
Inappropriate (composite)	Touch	2.82 (1.96)*	3.24 (2.00)*
	No	2.52 (1.80)	1.87 (1.60)
Surprising	Touch	3.96 (2.10)*	5.04 (1.61)*
	No	3.04 (2.10)	2.32 (2.03)

touch action made the robot appear more **Comforting** $F(1, 49) = 11.21$, $p = .002$, $d = 0.94$, more **Surprising**, $F(1, 49) = 28.07$, $p < .0001$, $d = 1.48$, and more **Inappropriate** (composite), $F(1, 49) = 7.4$, $p = .009$, $d = 0.76$. There was no effect on being **Measured** ($p = .54$, $d = 0.18$).

5.2.3 Comparison to Experiment I

Table 1 shows the means in Experiment II compared with the corresponding means in Experiment I. In both experiments, touch had a beneficial effect on trust, more so for Ethical & Sincere, but only in Experiment I did the effect reach significance. Although Experiment II tended to show stronger effects, the small sample size limited its statistical power. The effects of touch on behavior impressions (Comforting, Inappropriate composite, and Surprising) were all in the same direction as in Experiment I and strong enough to reach statistical significance even with modest sample size.

5.2.4 Multivariate Analyses

We aimed to replicate the mediation analyses from Experiment I, focusing on overall patterns and parameter estimates, because significance levels could not always be reached in this underpowered study. In the model of touch affecting trust via interpretations of Comforting, the initial impact of *Touch* on trust was positive ($b = 0.44$, $z = 1.26$, $p = .21$), but when introducing Comforting as a mediator, the direct effect turned negative ($b = -0.32$, $z = -1.05$, $p = .296$) while yielding a strong indirect effect ($b = 0.76$, $z = 2.94$, $p = .003$). This mediation pattern was once more stronger for Sincere & Ethical than for Capable & Reliable. Thus, the mediation results across the first two experiments are consistent, as illustrated in Fig. 4.

As in Experiment I, no other mediators reached traditional significance. Even when relaxing significance standards, no other variables showed consistent mediation trends across the two experiments. We did, however, replicate the finding from Experiment I that the set of behavior impressions (without the *Touch* manipulation) directly predicted trust, $R = .73, F(4, 46) = 13.0, p < .001$. Over and above Comforting ($r_{sp} = .32$), we found independent predictive contributions from Measured ($r_{sp} = .29$), Surprising ($r_{sp} = .27$), and the Inappropriate composite ($r_{sp} = -.22$), all $ts > 2.1$ and $ps < .012$.




Fig. 4 Mediation analyses of the positive effect of the *Touch* manipulation on overall trust, fully mediated by perceptions of the interpretation of the robot's behavior as being comforting. (Numbers are raw regression coefficients)

5.3 Discussion

Given the overall consistent results in the first two experiments, we can say that a robot touching a human elicits contrasting reactions. On the one hand, the robot is perceived as comforting and even tends to (indirectly) raise people's trust in it; on the other hand, it is seen as surprising and inappropriate. Experiment II attempted to reconcile this contrast by preceding the encounter with an initial handshake. But, in fact, the effect sizes on

all dependent variables were larger than in Experiment I, making the handshake seem, if anything, to backfire. The touch was seen as even more surprising and inappropriate in the context of the handshake in Experiment II than it was without such context in Experiment I (see Table 1). Handshakes have a formal and established social meaning [9], and in the context of such a formal relationship, touching the other's shoulder may have violated the implied formal relationship. Any benefit of the connection built by the handshake was then wiped out by the expectation violation; the net effect was again a robot that seemed to have both positive intentions but also to act inappropriately.

We should add that the No-Touch robot in Experiment II, which did engage in a physical handshake, was rated just as low on the descriptor Comforting as the No-Touch condition in Experiment I (which had no physical contact at all with the human). This suggests that the meaning of the shoulder touch in Experiment I was more specific than mere physical contact: people actually interpreted it as intending to comfort, whereas a handshake by itself was not interpreted that way.

One intriguing aspect of the mediation patterns in the first two studies was that the touch manipulation changed sign from positive (predicting increased trust) to negative (predicting decreased trust) when Comforting was included as a mediator. That is, the shoulder touch had a positive effect on perceptions of being comforting, and perceptions of being comforting had a positive effect on feelings of trust; but once we statistically controlled for this beneficial trust effect of Comforting, the shoulder touch lost its positive effect on trust. What was left was the (smaller) negative impression of touch as inappropriate, which is detrimental to trust. "Controlling for" is a difficult statistical concept, but we can capture it by saying that, if everybody had the same perceptions of the robot as being comforting (i.e., these varying perceptions are controlled for), then the shoulder touch, because it is also inappropriate, would be slightly decreasing trust.

In sum, participants in Experiment II fully recognized that the robot was intending to comfort the human, but they found that also surprising and inappropriate, especially when first framing the interaction as a formal relationship. Nonetheless, the power of perceiving the robot as comforting was strong enough to raise people's trust in the robot.

6 Experiment III

In Experiment III we continued our exploration of the context in which robot touch occurs and gets interpreted. Just like a handshake seems to serve as a social cue for a formal relationship, we investigated whether the robot’s *designed function* would serve as a cue for its proper interaction with a human. When a robot is specifically designed for a customer-service function (to optimize the experience of the customer with whom it interacts), we might expect a shoulder touch to be more appropriate (but still comforting) and Sincere & Ethical trust should benefit from that. When specifically designed for efficient performance, the robot may elicit Capable & Reliable trust, but there would be no expectation for a comforting gesture like a shoulder touch; this would be seen as inappropriate and lead to a drop in trust.

We used the standard video from Experiment I and II (female human actor and male robot with neutral attitude) to create a Baseline condition in Experiment III, and we introduced two new conditions, Customer-focused and Performance-focused design, which together make up the three-level *Function* factor. Moreover, in each of these conditions we presented half of participants with the handshake preceding the video, the other half without. And as before, the focal manipulation was the presence or absence of a shoulder touch.

6.1 Methods

6.1.1 Participants

Recruited through Amazon Mechanical Turk, 361 people participated in this study (female: 142, other: 3, no answer: 1) Their ages ranged from 18 to 69 ($M = 35.3$, $SD = 10.10$). The ethnic composition was as follows: 73.7% White or Caucasian, 9.3% Black or African American, 8.3% Asian, 7.0% Hispanic, 0.5% American Indian or Alaska Native, 0.3% Native Hawaiian or Pacific Islander, and 0.8% Other.

6.1.2 Procedure and Materials

The procedure was the same as in Experiments I and II, and we used the same combination of a female human actor, male robot, and neutral attitude. Reflecting the *Touch* condition, participants were randomly assigned to the version of the video with or without robot touch. Participants were also randomly assigned to a *Handshake* condition, which determined whether or not they saw

the handshake picture (Fig. 3) on the screen preceding the main video. Finally, in this experiment participants were randomly assigned to a *Function* condition, which determined which introductory text they saw before the video. The *Baseline* group had the same text on the introductory video screen as the first two experiments:

- **Baseline:** In the following scene you will see a human customer enter some information into a computer system (imagine entering your flight information when you check in at an airline counter). The robot assists the person. Please form an overall impression of the interaction.

The text for the two new designed function groups was different. Participants saw one of these introductory narratives:

- **Customer-Focused:** In the following scene you will see a human customer enter some information into a computer system (imagine entering your flight information when you check in at an airline counter). The robot assists the person. The robot is designed to be customer-focused and optimize the customer’s experience. Please form an overall impression of the interaction.
- **Performance-Focused:** In the following scene you will see a human customer enter some information into a computer system (imagine entering your flight information when you check in at an airline counter). The robot assists the person. The robot is designed to be performance-focused and to optimize the information-entering process. Please form an overall impression of the interaction.

The measures were the same as in Experiments I and II: the MDMT-based trust ratings and the six behavioral impression ratings.

6.2 Results

Twenty-six people (7.2%) failed the manipulation check question, leaving 335 participants for analysis. For each set of dependent variables we first analyzed the Baseline conditions, which presented the opportunity to replicate Experiments I and II. Then we examined the new conditions that manipulated designed function.

6.2.1 Trust

In the baseline group (without any design function information), trust did not vary as a function of the robot’s touch, either without handshake (corresponding to Experiment I), $F(1, 108) = 0.02$,

$p = .88$, $d = -0.04$, or with an initial handshake (corresponding to Experiment II), $F(1, 108) = 0.22$, $p = .64$, $d = -0.12$. These patterns were consistent for the two trust components, all F s < 1 and $ps > .60$ (see Table 2).

Table 2 Effects of robot touch on trust in Experiment III's baseline group (no function information) under conditions of No Handshake (like Experiment I) and With Handshake (like Experiment II)

		No Handshake	With Handshake
		Mean (SD)	Mean (SD)
All Trust	Touch	4.52 (1.44)	4.57 (1.82)
	No Touch	4.59 (1.62)	4.77 (1.57)
Sincere & Ethical	Touch	4.34 (1.66)	4.36 (1.92)
	No Touch	4.32 (1.94)	4.55 (1.64)
Capable & Reliable	Touch	4.70 (1.38)	4.78 (1.91)
	No Touch	4.86 (1.57)	4.99 (1.64)

Note: No comparisons between touch and no touch conditions (means on top of one another) were significant, $ps > .60$.

Next, we examined the *Touch* manipulation under conditions of designed function (Table 3). Overall, touch had a detrimental effect (unlike the beneficial effect we found in Experiments I and II): Participants trusted the touching robot less ($M = 4.06$, $SD = 1.80$, $n = 110$) than the non-touching robot ($M = 4.59$, $SD = 1.60$, $n = 114$), $F(1, 216) = 5.3$, $p = .022$, $d = -0.31$. However, this detrimental effect interacted with the presence of a handshake, $F(1, 216) = 3.8$, $p = .052$. Whereas loss of trust due to touch occurred in the no-handshake condition, $d = -0.59$, it was fully alleviated when people saw a robot-human handshake first, $d = 0.04$. This pattern was consistent across the two trust components and across the two design functions.

Table 3 Effects of robot touch on trust in Experiment III's groups with designed function information under conditions of No Handshake and With Handshake.

		No Handshake	With Handshake
		Mean (SD)	Mean (SD)
All Trust	Touch	3.73* (1.71)	4.42 (1.84)
	No Touch	4.67 (1.46)	4.50 (1.74)
Sincere & Ethical	Touch	3.52 (1.87)*	4.16 (2.06)
	No Touch	4.38 (1.65)	4.21 (2.01)
Capable & Reliable	Touch	3.93 (1.73)*	4.69 (1.72)
	No Touch	4.97 (1.47)	4.79 (1.60)

Note: Asterisk refers to a significant difference ($p = .05$) between

touch and no touch conditions (means on top of one another) for a given trust measure.

6.2.2 Behavior Impressions

The main results of the behavior impression analyses are illustrated in Fig. 5. We again formed an Inappropriate composite and analyzed the remaining items (Comforting, Surprising, Measured) on their own.

We first assessed behavior impressions as a function of touch, without or with handshake, solely in the Baseline group (the light and dark blue bars in Fig. 5). Without handshake (light blue bars; corresponding to Experiment I), the robot's touch had no significant impact on appearing Comforting, $F(1, 107) = 0.74$, $p = .39$, $d = 0.29$ (though the difference was in the familiar direction of more comforting intentions inferred from touch), while touch made the robot's behavior appear significantly more Surprising, $F(1, 107) = 8.4$, $p = .005$, $d = 0.80$, and more Inappropriate (composite), $F(1, 107) = 6.3$, $p = .013$, $d = -0.64$. There was no effect on perceptions of being Measured, $F(1, 107) = 0.31$, $p = .58$, $d = 0.16$. When seeing a handshake first (dark blue bars; corresponding to Experiment II), none of the behavior impressions differed significantly as a function of touch. Thus, without handshake we see a largely similar impression profile for touch vs. no touch as in Experiment I, but in the present participant sample, the handshake seemed to flatten this profile (unlike in Experiment II, where it maintained the profile).

Next, we examined the effects of *Touch* on behavior impressions under the new conditions of designed function (the light and dark red bars in Fig. 5). The two functions were indistinguishable, so we report only their averaged results. Without handshake (light red bars), the designed-function robot's touch had no impact on appearing Comforting, $F(1, 220) = 0.46$, $p = .50$, $d = 0.13$; but it made the robot appear more Surprising, $F(1, 220) = 13.3$, $p < .001$, $d = 0.68$, and substantially more Inappropriate (composite), $F(1, 220) = 30.4$, $p < .001$, $d = 1.04$. When seeing a handshake first, people tended to see the designed-function robot's touch as slightly more Comforting, $F(1, 220) = 3.4$, $p = .068$, $d = 0.33$, and they continued to see the robot as more Surprising, $F(1, 220) = 10.9$, $p = .001$, $d = 0.63$, more Inappropriate, $F(1, 220) = 10.6$, $p = .001$, $d = 0.62$, and also less Measured, $F(1, 220) = 10.5$, $p = .001$, $d = 0.65$. Overall, then, the designed-function robot's touch had mostly detrimental effects on behavior impressions, whether exposed to a robot-human handshake or not.

To confirm that our analysis strategy did not conceal possible differences between the design




Fig. 5 Behavior Impressions in Experiment III: Difference scores (with standard errors) of *Touch–No Touch*, within Baseline vs. Designed x No-Handshake Control vs. Handshake conditions.

function conditions and the baseline conditions, we reanalyzed the full design for all behavior impression variables, and the only effects that appeared were touch by handshake interactions that reflect the patterns described above. No main effect or interaction involving design function was close to significance.

6.2.3 Multivariate Analyses

The premises of the mediation analyses conducted in the previous two studies were not met in the present study: robot touch did not have a positive effect on trust, and it did not have a consistent impact on perceptions of being Comforting. In the new design function conditions, however, touch had a negative impact on trust and also a strong impact on the perceptions of the robot as being inappropriate and surprising. Therefore, we analyzed whether, in these design function conditions, perceptions of the robot as being surprising or inappropriate would mediate the negative impact of touch on trust. Surprising did not predict trust, so we limited the mediation analysis to the Inappropriate composite.

The initial impact of *Touch* on trust was negative ($b = -0.52, z = -2.31, p = .021$), as we had

seen earlier. Introducing Inappropriate as a mediator eliminated this direct effect ($b = -0.05, z = -0.19, p = .85$) while yielding a strong indirect effect ($b = -0.48, z = -3.65, p < .001$). As shown in Fig. 6, the touch manipulation increased ratings of Inappropriate by 2.1 points, and for every point increase in the Inappropriate ratings, trust decreased by 0.22 points.




Fig. 6 Mediation analysis of the detrimental effect of the *Touch* manipulation on overall trust, fully mediated by impressions of Inappropriateness (in designed function conditions of Experiment III).

Finally, we conducted a multiple linear regression of trust regressed on behavior impressions (ig-

noring the ineffective *Touch* manipulation). The results were very similar in all three Function conditions, so we report the data from the full sample. Prediction of trust from behavior impressions was substantial, $R = .70$, $F(4, 330) = 103.8$, $p < .001$. As in the first two studies, Comforting was the strongest predictor ($r_{sp} = .47$), and the other items made small but significant contributions: Measured ($r_{sp} = .18$), Inappropriate composite ($r_{sp} = -.17$), and Surprising ($r_{sp} = .09$), all $ts > 2.5$, $ps < .01$. These numbers are strikingly similar to those in Experiments I and II.

6.3 Discussion

In this experiment, we did not see the same effects of robot touch on trust that we saw in Experiments I and II. In those studies, a robot touching a human's shoulder increased an observer's trust in the robot by increasing perceptions of the robot being comforting (though touch also led to perceptions of Inappropriateness). In the comparable baseline condition of Experiment III, touch had nearly no impact on perceptions of being comforting and no impact on trust. In one important respect, this pattern is consistent with Experiments I and II: Because touch did not induce perceptions of the robot being comforting in Experiment III, it *therefore* did not increase trust in the robot (as it had in Experiments I and II). Perceptions of being comforting still strongly predicted trust in Experiment III; but the robot touch was simply not seen as particularly comforting overall (while it was still seen as more surprising and inappropriate, as in Experiments II and III).

In the two new designed function conditions, the robot touch became detrimental for trust. Statistically, that effect was entirely explained by increased perceptions of inappropriateness. The previous two experiments had already shown that robot touch has a double effect: it makes the robot appear both comforting and inappropriate. In Experiment III, the effect shifted predominantly to the inappropriate side. We can only speculate why information about designed function had this consequence: perhaps the awareness of a defined function made any behavior that did not fit the function's narrow expectation (whether performance-focused or service focused) seem deviant, resulting in a loss of trust. Even in the Baseline condition, the touch caused impressions of Inappropriate and Surprising; but only in the context of the designed functions did these impressions actually decrease trust. In that context, the touch perhaps made

people lose trust in the robot's ability to fulfill its designed function and not go beyond it.

7 General Discussion

Across all three studies, we showed that observing a robot's single touch gesture has a marked and consistent impact on an observer's impressions of the robot. Trust in the robot is also strongly and consistently influenced by impressions of the robot's behavior. But even though we have documented this consistent overall structure of robot touch, impressions, and trust, what varies across studies are the specific linkages between the three variables. People's specific interpretations of robot touch, and the conditions under which these interpretations increase or decrease trust, are not yet fully understood.

Nonetheless, two of the three studies showed that touch increases trust by way of increasing perceptions of comforting intentions (see Fig. 4. Even though this effect did not hold up in Experiment III, it is reasonable to ascribe this replication failure to sampling variations. Increasingly, psychological research has recognized that effects shown in one or more studies may not replicate in all studies. Sometimes the reasons are low statistical power, publication bias, p-hacking, or numerous other facets of the imperfect research endeavor [49]. But often it is simply sampling error [57]. The effect parameters in our third experiment are well within the confidence intervals of the first two; so we have no deep anomaly to puzzle over. We thus propose to retain, until contrary evidence comes to light, the hypothesis that there really is an effect of touch on trust, mediated by perception of being comforting.

If one accepts this provisional hypothesis, it should be noted that the psychological mediation of touch effects on trust may come with more than one facet. In addition to triggering a perception of the robot's comforting intentions, touch also triggers perceptions of inappropriateness. As long as the impression mix of *comforting+inappropriate* on balance favors the comforting side (as in Experiments I and II), trust increases. Once it favors the inappropriate side (as in Experiment III), trust can decline.

What are the conditions under which touch favors positive impressions and holds back negative impressions? And under what conditions does the opposite occur? We comment on two concrete variables we examined in the present studies and then speculate about two other potential moderators.

The first examined variable was an initial handshake between the robot and the human. Such a handshake appeared to somewhat worsen people’s behavior impressions in Experiment II, but it had the power to alleviate some negative impressions in the baseline group of Experiment III, and it actually alleviated the trust loss in the designed-function group of Experiment III. We cannot draw any confident conclusions from these results and suggest to other researchers to make the meaning of the handshake clearer (in our studies it was on a photograph and therefore open to interpretation).

The second variable we examined, namely presenting the robot as having a designed function, had a detrimental effect on people’s trust, which was, however, alleviated by a preceding handshake. The designed-function manipulation did not have a negative impact by itself; only when paired with the robot touch did it decrease trust and worsen behavior impressions. Both designed functions (customer service vs. performance) showed the same negative effects, so whatever the function manipulation did to people’s interpretations, it did it in a general way. We suspect that a robot with a specific designed function narrows people’s expectations for what the robot will do, and a surprising behavior such as a touch on the shoulder may be seen as deviating from these expectations. We know from considerable research in HRI that expectations matter a great deal, and they matter for trust not less than for other responses to robots [32].

One feature of robots that can set up powerful expectations is the robot’s appearance, especially its degree of humanlikeness. Systematic investigations are needed that span the full range of humanlikeness [39]. Such studies are almost certainly going to rely on an observational paradigm (like we did), because the number of distinct stimulus conditions and the number of required participants are nearly impossible to implement in laboratory research.

Another variable that may have strong effects on the complex relationship between touch and trust is the robot’s transparency. Our attempt to make the robot’s *function* transparent was only partially successful, but transparency has many shades. It may include the robot’s communication while it executes the touch (e.g., “I’m sorry” may highlight the intent to comfort), the observer’s understanding of the robot’s capacities (e.g., is the touch a controlled, intentional movement?), and information about the robot’s owner or programmer (and their motivations).

We might wonder about the larger impact of the finding that a robot’s simple touch on a hu-

man’s shoulder can, at least under some conditions, increase trust in the robot. As robots enter roles in which it is not only possible but expected for them to touch humans, such as healthcare robots, this finding offers some promise (for cases where trust in the robot is desirable). In emotionally significant situations like healthcare, the ethics of having robots assist or potentially replace humans in certain roles are precarious. Showing that people’s perceptions of how ethically trustworthy a robot is might be affected by something as minor as a robot touch highlights how delicate the perceptions of robots can be, and how much care and responsibility needs to go into robot interaction design. For one thing, there is the possibility that a gesture of touch makes a person over-trust a robot [19]. The robot’s act of touching a person may not provide any further information about the robot’s other capabilities or social skills; so it could prove problematic if a person infers such capabilities and skills based solely on observing a touch. For another, the touch could at times backfire, reducing trust that was perhaps built on expectations about the robot’s specific role. With the currently limited understanding of the (positive or negative) effects of robot touch, designers should be advised to stay away from adding such deeply human gestures to the machine’s action repertoire.

7.1 Limitations and Future Work

The present experiments used an observation-based paradigm, as a first step into exploring the many factors that go into the interpretations of robot touch. This is certainly a limitation, but it would be impossible to conduct such studies with 1000 or so participants as live in-lab experiments. Already in these observation studies, complexities emerged, and we believe that additional observation studies are needed to identify more clearly what the major impact factors are in people’s interpretation of, and trust in, robots that touch humans. But even when such clarity is reached in observation studies, we cannot be sure that actually being touched in real life will have the same effects. These studies with robots and participants as live interaction partners are still necessary. Yet, the results from observation studies can help to narrow the number of conditions that need to be investigated in lab-based interaction studies, thus making them feasible. Live interactions studies would also present greater opportunities to study trust behaviorally, rather than solely with self-report questionnaires.

Another limitation is that the task context was unchanged across the three studies and narrowly concerned an information-entering scenario. Future studies will certainly need to expand into other, emotionally more significant contexts. Future work should also extend the length of the scenarios, as our videos offered only a very brief human-robot interchange. A more full and varied presentation of the robot may change, or perhaps stabilize, the effects we have observed here.

A third limitation is that we did not collect cultural information about the participants involved in the study. Touch is a culturally interpreted form of non-verbal communication, which entails considerable individual differences in the way people make sense of a robot touching a person [41]. Future studies looking at robot touch should collect such information, along with more detailed verbal measures of people's concrete interpretations of the robot's behavior.

8 Conclusion

In this paper, we set out to explore the impact of observing robot-initiated touch on human-robot trust. In two experiments we found that observing touch increased trust, but a third study did not replicate this pattern and hinted at conditions under which such touch can even decrease trust. Perceptions of the touch gesture as comforting were largely consistent, and so was the striking concurrent perception of the robot as being inappropriate. Additional context information, such as a handshake or design information, had inconsistent impact. Just as the touch itself appears to be open to multiple interpretations, so is the context in which the touch occurs. This indeterminacy reminds us of how enormously complex people's perceptions of robots are, even just in short video observations.

People still have very few experiences with robots, and in interpreting the experiences they do have, they must rely on the psychological mechanisms and habits that have evolved, biologically and culturally, from interactions with other humans. No new adaptations for robots exist, or at least not yet. Now that robots are demanding increasing attention and interaction, these psychological mechanisms may either have to be re-tuned or robot design needs to use all the knowledge we can muster to adapt to the complex human responses. It is clear that our knowledge is currently at a very early stage.

Observing small social signals like handshake or social touch can have a profound effect on the

perception of human-robot interactions and can significantly impact trust perceptions, everything else being equal. These effects can be further modulated based on professional and social roles, and it is thus important to not investigate these effects in isolation of background and task-based aspects. Moreover, observing different types of touch (socially sanctioned like handshake and typically non-sanctioned, possibly less appropriate shoulder touch) can interact in complex ways and thus deserve detailed follow-up to investigate how sanctioned and non-sanctioned touch can affect user perception and trust.

Funding: This project was funded in part by National Science Foundation grant IIS-1316809 and robot AFOSR grant FA9550-18-1-0465.

Conflict of Interest: The authors declare that they have no conflict of interest.

References

1. P. A. Andersen and K. Leibowitz. The development and nature of the construct touch avoidance. *Environmental psychology and nonverbal behavior*, 3(2):89–106, 1978.
2. R. Andreasson, B. Alenljung, E. Billing, and R. Lowe. Affective touch in human-robot interaction: conveying emotion to the nao robot. *International Journal of Social Robotics*, 10(4):473–491, 2018.
3. T. Arnold and M. Scheutz. 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*, pages 352–360. ACM, 2018.
4. J. Avelino, F. Correia, J. Catarino, P. Ribeiro, P. Moreno, A. Bernardino, and A. Paiva. The power of a hand-shake in human-robot interactions. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1864–1869. IEEE, 2018.
5. J. Bezemer and G. Kress. Touch: A resource for making meaning. *Australian Journal of Language & Literacy*, 37(2):77–85, June 2014.
6. S.-J. Blakemore, D. Bristow, G. Bird, C. Frith, and J. Ward. Somatosensory activations during the observation of touch and a case of vision-touch synaesthesia. *Brain: A Journal of Neurology*, 128(Pt 7):1571–1583, July 2005.
7. N. Bock, L. Hoffmann, and A. Rosenthal-vd Pütten. Your touch leaves me cold, robot. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 71–72. ACM, 2018.
8. BodyLanguageCentral. What it means when someone touches your shoulder. Retrieved from <https://bodylanguagecentral.com/what-it-means-when-someone-touches-your-shoulder/>, Jan. 2019.
9. J. K. Burgoon. Relational message interpretations of touch, conversational distance, and posture. *Journal of Nonverbal behavior*, 15(4):233–259, 1991.
10. A. W.-Y. Chan and C. I. Baker. Seeing is not feeling: Posterior parietal but not somatosensory cortex engagement during touch observation. *Journal of Neuroscience*, 35(4):1468–1480, Jan. 2015.

11. T. L. Chen, C.-H. King, A. L. Thomaz, and C. C. Kemp. Touched by a robot: An investigation of subjective responses to robot-initiated touch. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 457–464. IEEE, 2011.
12. C. Crowell, M. Scheutz, P. Schermerhorn, and M. Villano. Gendered voice and robot entities: Perceptions and reactions of male and female subjects. In *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, St. Louis, MO, October 2009.
13. A. H. Crusco and C. G. Wetzel. The midas touch: The effects of interpersonal touch on restaurant tipping. *Personality and Social Psychology Bulletin*, 10(4):512–517, 1984.
14. E. G. Dougherty and H. Scharfe. Initial formation of trust: designing an interaction with humanoid-robot to promote a positive attitude for cooperation. In *International Conference on Social Robotics*, pages 95–103. Springer, 2011.
15. D. Erceau and N. Guéguen. Tactile contact and evaluation of the toucher. *The Journal of social psychology*, 147(4):441–444, 2007.
16. J. D. Fisher, M. Rytting, and R. Heslin. Hands touching hands: Affective and evaluative effects of an interpersonal touch. *Sociometry*, 39(4):416–421, 1976.
17. H. Fukuda, M. Shiomi, K. Nakagawa, and K. Ueda. 'midas touch' in human-robot interaction: Evidence from event-related potentials during the ultimatum game. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, pages 131–132. ACM, 2012.
18. J. A. Hall. Touch, status, and gender at professional meetings. *Journal of Nonverbal Behavior*, 20(1):23–44, 1996.
19. P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. Chen, E. J. De Visser, and R. Parasuraman. A meta-analysis of factors affecting trust in human-robot interaction. *Human factors*, 53(5):517–527, 2011.
20. T. Hirano, M. Shiomi, T. Iio, M. Kimoto, I. Tanev, K. Shimohara, and N. Hagita. How do communication cues change impressions of human-robot touch interaction? *International Journal of Social Robotics*, 10(1):21–31, 2018.
21. K. A. Hoff and M. Bashir. Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3):407–434, May 2015.
22. L. Hoffmann. *That Robot Touch that Means so Much: On the Psychological Effects of Human-Robot Touch*. PhD thesis, University of Duisburg-Essen, Germany, 2017.
23. J. Hornik. Tactile stimulation and consumer response. *Journal of Consumer Research*, 19(3):449–458, 1992.
24. L. Hung, C. Liu, E. Woldum, A. Au-Yeung, A. Berndt, C. Wallsworth, N. Horne, M. Gregorio, J. Mann, and H. Chaudhury. The benefits of and barriers to using a social robot PARO in care settings: a scoping review. *BMC Geriatrics*, 19(1):232, Aug. 2019.
25. JASP Team. JASP (Version 0.11.1)[Computer software], 2019.
26. S. E. Jones and A. E. Yarbrough. A naturalistic study of the meanings of touch. *Communications Monographs*, 52(1):19–56, 1985.
27. S. M. Jourard. An exploratory study of body-accessibility. *British Journal of Social and Clinical Psychology*, 5(3):221–231, 1966.
28. T. Law, M. Chita-Tegmark, and M. Scheutz. The interplay between emotional intelligence, trust, and gender in human-robot interaction. *International Journal of Social Robotics*, pages 1–13, 2020.
29. T. Law and M. Scheutz. Trust: Recent concepts and evaluations in human-robot interaction. In C. S. Nam and J. B. Lyons, editors, *Trust in human-robot interaction: Research and applications*, chapter 2. Elsevier, San Diego, CA, 2021.
30. J. D. Lee and K. A. See. Trust in automation: Designing for appropriate reliance. *Human factors*, 46(1):50–80, 2004.
31. J. Levav and J. J. Argo. Physical contact and financial risk taking. *Psychological Science*, 21(6):804–810, June 2010.
32. B. F. Malle, K. Fischer, J. E. Young, A. Moon, and E. C. Collins. Trust and the discrepancy between expectations and actual capabilities of social robots. In D. Zhang and B. Wei, editors, *Human-robot interaction: Control, analysis, and design*. Cambridge Scholars Publishing, New York, NY, in press.
33. B. F. Malle and D. Ullman. A multi-dimensional conception and measure of human-robot trust. In C. S. Nam and J. B. Lyons, editors, *Trust in human-robot interaction: Research and applications*, chapter 1, pages 3–25. Elsevier, San Diego, CA, 2021.
34. K. Nakagawa, M. Shiomi, K. Shinozawa, R. Matsumura, H. Ishiguro, and N. Hagita. Effect of robot's active touch on people's motivation. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 465–472. ACM, 2011.
35. C. I. Nass and S. Brave. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, MA, 2005.
36. J. Nie, M. Park, A. L. Marin, and S. S. Sundar. Can you hold my hand? physical warmth in human-robot interaction. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 201–202. IEEE, 2012.
37. T. Nomura. Robots and gender. In M. J. Legato, editor, *Principles of Gender-Specific Medicine: Gender in the Genomic Era*, pages 695–703. Elsevier, 2017.
38. E. Park and J. Lee. I am a warm robot: the effects of temperature in physical human-robot interaction. *Robotica*, 32(1):133–142, 2014.
39. E. Phillips, X. Zhao, D. Ullman, and B. F. Malle. What is human-like? Decomposing robots' human-like appearance using the Anthropomorphic roBOT (ABOT) database. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18*, pages 105–113. ACM, New York, NY, USA, 2018.
40. A. Powers, A. D. Kramer, S. Lim, J. Kuo, S.-I. Lee, and S. Kiesler. Eliciting information from people with a gendered humanoid robot. In *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, pages 158–163. IEEE, 2005.
41. M. S. Remland, T. S. Jones, and H. Brinkman. Interpersonal distance, body orientation, and touch: Effects of culture, gender, and age. *The Journal of social psychology*, 135(3):281–297, 1995.
42. S. Šabanović, C. C. Bennett, W.-L. Chang, and L. Huber. Paro robot affects diverse interaction modalities in group sensory therapy for older adults with dementia. In *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*, pages 1–6. IEEE, 2013.

43. P. Schermerhorn, M. Scheutz, and C. R. Crowell. Robot social presence and gender: Do females view robots differently than males? In *Proceedings of the Third ACM IEEE International Conference on Human-Robot Interaction*, pages 263–270, Amsterdam, The Netherlands, 2008. ACM Press.
44. T. Shibata, T. Mitsui, K. Wada, A. Touda, T. Kumasaka, K. Tagami, and K. Tanie. Mental commit robot and its application to therapy of children. In *2001 IEEE/ASME International Conference on Advanced Intelligent Mechatronics. Proceedings (Cat. No. 01TH8556)*, volume 2, pages 1053–1058. IEEE, 2001.
45. T. Shibata and K. Tanie. Physical and affective interaction between human and mental commit robot. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*, volume 3, pages 2572–2577. IEEE, 2001.
46. T. Shibata, T. Tashima, and K. Tanie. Subjective interpretation of emotional behavior through physical interaction between human and robot. In *IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on Systems, Man, and Cybernetics (Cat. No. 99CH37028)*, volume 2, pages 1024–1029. IEEE, 1999.
47. M. Shiomi, K. Nakagawa, K. Shinozawa, R. Matsumura, H. Ishiguro, and N. Hagita. Does a robot's touch encourage human effort? *International Journal of Social Robotics*, 9(1):5–15, 2017.
48. M. Shiomi, A. Nakata, M. Kanbara, and N. Hagita. A hug from a robot encourages prosocial behavior. In *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*, pages 418–423. IEEE, 2017.
49. J. P. Simmons, L. D. Nelson, and U. Simonsohn. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11):1359–1366, Nov. 2011.
50. M. Strait, C. Canning, and M. Scheutz. Let me tell you! investigating the effects of robot communication strategies in advice-giving situations based on robot appearance, interaction modality and distance. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 479–486, 2014.
51. C. Torrey, S. R. Fussell, and S. Kiesler. How a robot should give advice. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 275–282. IEEE, 2013.
52. D. Ullman and B. F. Malle. The Multidimensional Measure of Trust (MDMT), v1. Retrieved from http://research.clps.brown.edu/SocCogSci/Measures/MDMT_v1.pdf.
53. D. Ullman and B. F. Malle. What does it mean to trust a robot?: Steps toward a multidimensional measure of trust. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 263–264. ACM, 2018.
54. D. Ullman and B. F. Malle. Measuring gains and losses in human-robot trust: Evidence for differentiable components of trust. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 618–619. IEEE, 2019.
55. C. J. Willemse, A. Toet, and J. B. van Erp. Affective and behavioral responses to robot-initiated social touch: toward understanding the opportunities and limitations of physical contact in human-robot interaction. *Frontiers in ICT*, 4:12, 2017.
56. C. J. Willemse and J. B. van Erp. Social touch in human-robot interaction: Robot-initiated touches can induce positive responses without extensive prior bonding. *International journal of social robotics*, 11(2):285–304, 2019.
57. B. M. Wilson, C. R. Harris, and J. T. Wixted. Science is not a signal detection problem. *Proceedings of the National Academy of Sciences*, 117(11):5559–5567, Mar. 2020.