



I Can See It in Your Eyes: Gaze as an Implicit Cue of Uncanniness and Task Performance in Repeated Interactions With Robots

Giulia Perugia^{1*}, Maike Paetzel-Prüsmann², Madelene Alanenpää¹ and Ginevra Castellano¹

¹ Uppsala Social Robotics Lab, Department of Information Technology, Uppsala University, Uppsala, Sweden,

² Computational Linguistics, Department of Linguistics, University of Potsdam, Potsdam, Germany

OPEN ACCESS

Edited by:

German I. Parisi,
University of Hamburg, Germany

Reviewed by:

Markus Vincze,
Vienna University of Technology,
Austria
Hideyuki Nakanishi,
Osaka University, Japan

*Correspondence:

Giulia Perugia
giulia.perugia@it.uu.se

Specialty section:

This article was submitted to
Human-Robot Interaction,
a section of the journal
Frontiers in Robotics and AI

Received: 24 December 2020

Accepted: 04 March 2021

Published: 07 April 2021

Citation:

Perugia G, Paetzel-Prüsmann M,
Alanenpää M and Castellano G (2021)
I Can See It in Your Eyes: Gaze as an
Implicit Cue of Uncanniness and Task
Performance in Repeated Interactions
With Robots.
Front. Robot. AI 8:645956.
doi: 10.3389/frobt.2021.645956

Over the past years, extensive research has been dedicated to developing robust platforms and data-driven dialog models to support long-term human-robot interactions. However, little is known about how people's perception of robots and engagement with them develop over time and how these can be accurately assessed through implicit and continuous measurement techniques. In this paper, we explore this by involving participants in three interaction sessions with multiple days of zero exposure in between. Each session consists of a joint task with a robot as well as two short social chats with it before and after the task. We measure participants' gaze patterns with a wearable eye-tracker and gauge their perception of the robot and engagement with it and the joint task using questionnaires. Results disclose that aversion of gaze in a social chat is an indicator of a robot's uncanniness and that the more people gaze at the robot in a joint task, the worse they perform. In contrast with most HRI literature, our results show that gaze toward an object of shared attention, rather than gaze toward a robotic partner, is the most meaningful predictor of engagement in a joint task. Furthermore, the analyses of gaze patterns in repeated interactions disclose that people's mutual gaze in a social chat develops congruently with their perceptions of the robot over time. These are key findings for the HRI community as they entail that gaze behavior can be used as an implicit measure of people's perception of robots in a social chat and of their engagement and task performance in a joint task.

Keywords: perception of robots, long-term interaction, mutual gaze, engagement, uncanny valley

1. INTRODUCTION

An essential precondition for understanding the development of people's perception of robots in repeated interactions is the development of measurement techniques suitable for long-term assessment. To date, the measurement of people's perception of robots relies almost solely on questionnaires and interviews. However, these have several limitations. First, they only capture people's perception at one specific moment in time. This means that while changes in perception can be detected between the different points of measurement, it is not possible to relate these changes to particular events within the interaction. Second, in order to capture changes in people's perception

over time as accurately as possible, multiple points of measurement are required. However, filling out questionnaires interrupts people's interactive experience and hence has a high potential for decreasing the involvement with the robot and the task they perform with it. Finally, measures of self-report are prone to bias. Repeatedly filling out the same questionnaire may cause people to remember previous answers, which can decrease the accuracy of the measurement due to response fatigue, and reveal the purpose of the experiment due to learning and hypothesis guessing (Choi and Pak, 2005). To accurately study people's perception of robots in repeated interactions and relate changes in perception to specific actions of the robot, it is thus important to develop more implicit and continuous measurement techniques. In this paper, we explore *gaze patterns as a potential continuous method for capturing people's perception of robots*.

Previous studies in Social Psychology investigating the relationship between gaze behavior and people's perception of a human partner came to conflicting results. Indeed, people gaze more at each other when they share feelings of warmth and liking and when seeking friendship (Kleinke, 1986; Wirth et al., 2010; Cui et al., 2019). However, they show similar behavior patterns (i.e., longer fixation) when interacting with unconventionally-looking people, for instance, those carrying a facial stigma (Madera and Hebl, 2012, 2019). This is because a stimulus that does not match prior knowledge and expectations captures more attention than a stimulus that perfectly matches them (Langer et al., 1976; Bless and Greifeneder, 2017). Transferring this to human-robot interaction (HRI), we can expect people to look more often to robots that they like, as well as to stare longer at unconventional robots, such as realistic androids eliciting uncanny feelings (Mori et al., 2012). Indeed, when presented with uncanny robots, people have sometimes avoided looking at them (Strait et al., 2015) and other times stared at them longer (Minato et al., 2004; Thepsonthorn et al., 2021). In social robotics, experiments on the meaning of mutual gaze have almost solely focused on uncanny robots and still images. In this paper, we hence aim to *discover whether findings from non-interactive scenarios translate to real face-to-face interactions with robots and whether mutual gaze in a social chat can be an implicit measure of people's perception of robots, in particular of likability and uncanniness*.

In the literature, the few experiments that tracked gaze toward a robot in actual interactions between a human and a robot mainly used joint tasks involving multiple objects of attention (e.g., touchscreen) as a test-bed (e.g., Castellano et al., 2009; Kennedy et al., 2015; Papadopoulos et al., 2016). We argue that in these contexts, the gaze toward the robot has a meaning distinct from the one it has in a face-to-face social conversation, as the robot and the objects involved in the joint task compete for the same attentional resources. The second focus of this paper is thus to *understand whether the gaze participants allocate to the robot in a joint task is related to engagement and task performance and what is the meaning of the gaze people direct to the other objects involved in the joint task*. Previous long-term HRI studies on gaze exclusively focused on how the gaze toward the robot developed over multiple sessions of the same activity

(e.g., Serholt and Barendregt, 2016; Ahmad et al., 2017; Ahmad and Mubin, 2018). In this experiment, *we also investigate how the gaze toward other foci of attention in the joint task varies over time and whether participants' mutual gaze in a social chat preceding and succeeding the joint task changes across repeated interaction sessions*.

This paper presents an exploratory study in which participants were involved in three interaction sessions with the blended robotic head Furhat (Al Moubayed et al., 2012) occurring with multiple days of zero exposure in between. To achieve meaningful variations in the robot's perception, we manipulated its humanlikeness by applying three facial textures with different anthropomorphic features. Each interactive session was divided into a geography-themed cooperative game with the robot (serving as a joint task) and a face-to-face social chat before and after the game. To track and analyze gaze patterns, participants wore eye-tracking glasses throughout the interactive session. At different points in the sessions, they were asked to self-report their perception of the robot and their engagement with it and the collaborative game. The questionnaires were used to gain novel insights into the suitability of gaze patterns as an implicit measure of people's perception of a robot and of their engagement and performance in a joint task. The multiple sessions of interaction enabled us to track the progression of gaze within and between interactions and understand if and how gaze patterns change over time.

2. RELATED WORK

2.1. Mutual Gaze and Liking

Goffman (1964) was one of the first to state that the direction of gaze plays a crucial role in the initiation and maintenance of social encounters and can be an indicator of social attention. Exline et al. (1965) showed how the amount of mutual gaze increases when a person is drawn by another individual, either in an affiliative or competitive way. It is through the mutually held gaze that two people commonly establish their openness to another's communication, and the aversion of the eyes in a face-to-face interaction can be read as a cut-off act as well as a sign of dislike (Kendon, 1967; Ray and Floyd, 2006).

Studies on mutual gaze in HRI mostly focused on how the implementation of such nonverbal behavior on a robot influences users' perception (Mumm and Mutlu, 2011; Kompatsiari et al., 2017). A 2017 review identified three main lines of gaze research in HRI: (1) human responses to robot's gaze, (2) design of gaze features for robots, and (3) computational tools to implement social gaze in robots (Admoni and Scassellati, 2017). The first attempt to use gaze as a mean to assess interest, liking, and engagement in HRI was made by Sidner et al. (2005) who involved participants in a demo interaction with the penguin robot Mel and used gaze to understand whether the manipulation of the robot's behavior influenced the amount of mutual gaze it attracted. Lemaignan et al. (2016) measured the direction of gaze of children involved in a collaborative task with the NAO robot (i.e., teaching handwriting skills to a robot) and used it to compute their *with-me-ness*, the extent a human is *with* the robot over the course of an interactive task. They

obtained a with-me-ness value by comparing the child's focus of attention at a certain point in time with a set of expected attentional targets for that moment of the interaction. Kennedy et al. (2015) found children interacting with a physical robot to gaze significantly more often to the robot than children interacting with a virtual one, and that they spent significantly more seconds per minutes gazing at the robot in the real robot condition than in the virtual one. Similarly, Papadopoulos et al. (2016) used gaze to estimate the social engagement with the robot of adults in a memory game with the NAO robot, and Castellano et al. (2009, 2010) of children in a chess game with the iCat robot.

These related studies seem to suggest that the most compelling interaction conditions are those that elicit the longest gaze toward the robot, partially supporting the positive mutual gaze-liking relationship in the context of HRI. However, only a few of these studies specifically related participants' gaze toward the robot with metrics of likability and used it as an implicit measure of participants' perception of a robot (Sidner et al., 2005). Moreover, most of the reviewed studies focused on the allocation of gaze toward a robot during a task (e.g., a chess game). In such a context, the robot and the task at hand compete for the same attentional resources, hence gaze toward the robot is not anymore a precise measure of the robot's likability and of participants' social syntony with it because it is hindered by participants' willingness to complete the task (Corrigan et al., 2013, 2015; Perugia et al., 2018, 2020). In this study, *we thus examine robot-directed gaze in two separate situations: during a collaborative game, but also in a face-to-face social chat between the participant and the robot occurring before and after the game interaction.* In the former, we focus on the mutual gaze that the robot attracts and use it as a predictor of participants' perceptions. In the latter case, we focus on participants' gaze patterns and examine whether these can predict task performance, perceived involvement with the robot, and with the game. This is with the aim *to understand whether mutual gaze in a face-to-face social chat increases with the robot's likability and which gaze patterns are related with task performance and engagement in the joint task.*

2.2. Stigma, Staring, and the Uncanny Valley

Staring is defined as gaze that persists regardless of the behavior of the other person (Kleinke, 1986). The novel stimulus hypothesis posits that behavioral avoidance of people that appear as physically different (e.g., pregnant) is mediated by a conflict over a desire to stare at novel stimuli and a desire to adhere to a norm against staring when the novel stimulus is another person (Langer et al., 1976). Langer and colleagues discovered that, when staring is not negatively sanctioned, it varies as a function of the novelty of the observed subject, whereas, when the norms against it are instated, staring is inhibited. In line with this, Kleck (1968) found out that participants looked at a research confederate carrying a physical stigma significantly more than at one not carrying it and Madera and Hebl (2012) found that interviewers of facially stigmatized interviewees (i.e., port stain) spent considerably more

time looking at the specific location of applicants' stigma than interviewers evaluating non-stigmatized applicants.

In HRI, it is known that the likability of a robot increases with its humanlikeness up to a point where it drops abruptly. This drop in likability, known as the uncanny valley, is reached when a robot is almost indistinguishable from a healthy human, but some of its features still point to its artificiality and hence elicit eeriness (Mori et al., 2012). In their 2015 review on research related to the uncanny valley, Kätsyri et al. (2015) found extensive empirical evidence for the existence of the uncanny valley effect in at least some humanlike robots and outlined two competing explanatory theories behind the effect. On the one hand, the *perceptual mismatch theory* states that any conflicting cues in an agent's appearance can lead to uncanny feelings. On the other hand, the *categorical ambiguity theory* claims that only robots with conflicting cues leading to uncertainty about their categorical affiliation lead to uncanny feelings. As uncanny robots often feature atypical cues in their appearance, they might be perceived as more novel and the eeriness they generate might be equated to that elicited by a stigma. In this sense, one can hypothesize that robots perceived as uncanny elicit higher staring than robots that are not perceived as such. However, in line with the extant literature on the positive relationship between liking and mutual gaze, one can also posit that uncanny robots attract less direct gaze, as they are less likable and elicit more discomfort.

Minato et al. (2004) were the first to investigate whether the uncanniness of an android robot could have an effect on people's gaze behavior. They gauged the direction of gaze of people involved in a face-to-face conversation with three interlocutors: a human girl, a motionless android robot shaped as a girl, and the same android robot with a moving head, eyes, and neck. They found people to look significantly more at the eyes of the android robots than at those of the human girl and consequently suggested fixation time to be an implicit measure of uncanniness. Strait et al. (2015) exposed participants to pictures of real humans and robots varying in humanlikeness (low, medium, high) and found participants to fixate highly humanlike robots less than the other agents when the whole body was taken into account, and more than the artificial agents when the head and the eyes were considered. Smith and Wiese (2016) studied the effects of a robot's appearance on delayed disengagement. They asked participants to orient their gaze to a target dot appearing on the sides of a screen after fixating an agent in the center of it and measured the time it took for participants to reorient their gaze. Although reaction times should increase when processing stimuli with a negative connotation, their results did not disclose any significant difference across agents varying in humanlikeness (e.g., non-social, robot, robotoid, humanoid, human). A similar study was carried out by Li et al. (2015), who investigated both static and video stimuli. In the static image experiment, the reaction times to the mechanical robot were slower than those to the android robot and real human. In the video-based experiment, on the contrary, the reaction times to the android robot and the real human were slower than those to the mechanical robot. Since these related studies are mostly focused

on non-interactive stimuli and their results do not point to a clear direction with respect to the two alternative hypotheses on the meaning of mutual gaze, we further explore whether *mutual gaze in a face-to-face interaction with a social robot varying in humanlikeness is related to its perceived uncanniness, and, if so, whether this relation aligns with the mutual gaze-liking or novel stimulus hypothesis.*

2.3. Tracking Gaze Over Time

When it comes to the progression of gaze over time in interactive scenarios with social robots, most of the related work is focused on Child-Robot Interaction (cHRI). In this context, pivotal work has been performed by Baxter et al. (2014) and Kennedy et al. (2015) who focused on changes in gaze patterns within an interaction session. Baxter and colleagues measured children's gaze behavior toward the robot during a joint task by calculating a number of gaze metrics (i.e., mean length of gaze to the robot and length of gaze to the robot per minute) within a predefined time-window and comparing them with the same metrics gauged in subsequent time-windows. By splitting the interactions in three equal parts, they found that the gaze directed to the robot decreased from the first to the final third of the joint task and interpreted this result as a decrease in the engagement with the robot over time. Kennedy et al. (2015) used the same approach in a collaborative sorting task and noticed that the gaze toward the robot significantly reduced between the first and the second third of the interaction and then stayed more or less constant. Similar to Baxter et al., they ascribed this drop and subsequent stabilization to a reduction of engagement over time due to the wearing-off of the novelty effect.

Along this line, but with a stronger focus on long-term cHRI is the work of Serholt and Barendregt (2016). They involved 30 children in three sessions of play with a NAO robot in a map reading task and analyzed children's behavioral reactions to three implicit probes: a greeting, a feedback/praise, and a question. One of the behavioral markers employed to assess children's reactions to the probes was the gaze toward the robot, which they considered a sign of social engagement. Serholt and Barendregt found that the most common response to the three probes was directing the gaze toward the robot, and that over time, this response decreased slightly. The authors suggested that one way to counteract this decrease in children's engagement with the robot over time was to implement responsive robot behaviors that could facilitate bonding. Ahmad and colleagues moved in this direction by studying how different types of robot adaptation to children's states could influence social engagement and learning (Ahmad et al., 2017, 2019; Ahmad and Mubin, 2018). They ran several long-term studies (three to four sessions) involving children in joint tasks with a NAO robot (i.e., snakes and ladders game, mathematical learning task, vocabulary learning task) and evaluated the effect of different types of robot's adaptation (e.g., memory and emotion adaptation) on children's engagement with it. They measured children's social engagement with the robot through a number of behavioral metrics, among which the gaze directed to the robot. As postulated by Serholt and Barendregt (2016), they

found that the gaze allocated to the robot during the joint task increased across sessions when the robot behaved empathetically (Ahmad et al., 2017; Ahmad and Mubin, 2018) and that children learned significantly more over time when interacting with the empathetic robot (Ahmad and Mubin, 2018) or when the robot gave them positive and supportive feedback (Ahmad et al., 2019).

The literature discussed above shows that gaze has been consistently used to measure social engagement with robots over time. However, in most cases, gaze has been manually annotated (Baxter et al., 2014; Kennedy et al., 2015; Serholt and Barendregt, 2016; Ahmad et al., 2017, 2019; Ahmad and Mubin, 2018). While several researchers have proposed automated methods to gauge gaze allocation (Anzalone et al., 2015; Lemaignan et al., 2016; Lala et al., 2017), only Del Duchetto et al. (2020) have used such methods to assess the development of social engagement with robots over the time of an interaction, and, to the best of our knowledge, no one has used them to monitor the direction of gaze toward different foci across repeated interactions. For this study, *we automatically annotate gaze with a deep learning-based object detection algorithm utilizing YOLOv4 (Bochkovskiy et al., 2020), and investigate how gaze patterns in a joint task develop between three interaction sessions with multiple days of zero exposure in between.* We believe that automatic gaze tracking holds promises for online assessment of engagement and could be used for real-time reward estimation in co-adaptive scenarios in the future.

The main focus of long-term gaze studies has been engagement. However, Strait et al. (2015) and Minato et al. (2004) show how gaze can also be a meaningful predictor of a robot's uncanniness. From our previous work, we know that: (1) the mere exposure to a robot changes people's initial perceptions of it; (2) progressively exposing people to the multimodal behaviors of a robot improves people's perception of it (Paetzel and Castellano, 2019); and (3) the perceptual dimensions that contribute to people's mental image of the robot stabilize over time (Paetzel et al., 2020; Paetzel-Prüsmann et al., 2021). Hence, besides investigating the role of gaze patterns in a joint task with a focus on engagement, in this paper, we also focus on *understanding how mutual gaze in a social chat develops over time within and between interaction sessions and how it relates to people's perception of the robotic interaction partner.* Indeed, if mutual gaze was found to be a meaningful predictor of people's perception of robots, it could be used to track the development of people's mental image of a robot over time. To the best of our knowledge, this approach has never been attempted before.

3. RESEARCH QUESTIONS

This exploratory work aims to further our understanding of the meaning of gaze in two types of interactions with robots: face-to-face social chats and a joint task. In the former, we focus on mutual gaze and attempt to understand whether it is related to people's *perception* of the robot. In the latter, we focus on people's gaze toward the robot and other objects involved in the game and explore which gaze pattern is related to participants' *task*

performance, involvement with the game, and involvement with the robot. Hence, we pose the following research questions:

- RQ1** Is the mutual gaze directed to the robot in a face-to-face social chat a predictor of people's perception of the robot?
RQ2 Which gaze pattern in a joint task is predictive of people's engagement and task performance?

Extant literature has found that gaze toward a robot decreases over the time of an interaction (Minato et al., 2004; Baxter et al., 2014; Kennedy et al., 2015). Similarly, we attempt to understand whether mutual gaze toward the robot reduces between two equally long *social chats* occurring before and after a joint task. Moreover, we explore whether it changes over three repeated interaction sessions. This way, we aim to answer the following research question:

- RQ3** Does the mutual gaze directed to the robot change between a pre- and post-game face-to-face social chat and across repeated interactions?

Previous research has further discovered that the amount of gaze directed to a robot in a *joint task* slightly declines across repeated interactions (Serholt and Barendregt, 2016; Ahmad et al., 2017; Ahmad and Mubin, 2018). As these works have overlooked the gaze participants direct to other objects involved in the joint task (e.g., tablet and touchscreen), it is difficult to establish whether the decline in the gaze toward the robot they observe really corresponds to the allocation of attentional resources elsewhere. In this paper, we gauge both the gaze directed to the robot and the gaze directed to the other objects involved in the game (e.g., tablet and touchscreen) and attempt to understand how gaze as a whole changes over repeated interactions. Thus, we pose the following research question:

- RQ4** Do gaze patterns in a joint task change across repeated interactions?

Since it has been shown that a robot's level of humanlikeness affects the amount of gaze it attracts (Minato et al., 2004; Strait et al., 2015), in this study, we vary the humanlikeness of the robot with which participants interact. This way, we aim to answer the following research questions:

- RQ5a** Does the level of humanlikeness of the robot affect the amount of mutual gaze directed to it in a face-to-face social chat?
RQ5b Does the level of humanlikeness of the robot affect people's gaze patterns during the joint task?

With respect to previous research which mainly focused on android robots and compared them with less humanlike robotic platforms (Minato et al., 2004; Strait et al., 2015), we keep the robot's embodiment constant across conditions by using a blended embodiment, and manipulate the humanlikeness of the robot exclusively by changing its facial texture.

4. METHODOLOGY

We designed an experiment involving participants in three *interaction sessions* (within-subject variable) with a social

robot displaying three levels of *humanlikeness* (between-subject variable): humanlike, mechanical, and a morph between the two (cf. **Figure 1**). The interaction sessions had an average of 6.9 days of zero exposure in between (S1–S2: $M = 6.76$, $SD = 1.83$; S2–S3: $M = 7.05$, $SD = 2.41$). Each session was divided into three phases: (1) a social chat with the robot, (2) a joint task to perform, and (3) a final social chat.

4.1. Participants

As we suspected strong effects for the study, an initial check using G*Power (MANOVA: Repeated measures, within-between interaction, $\alpha = 0.05$, number of groups = 3, number of measurements = 3), considering strong effects $f(V) = 0.4$, resulted in a sample size of 61 participants. Hence, we recruited 60 participants from an international Master's course in Computer Science at Uppsala University to participate in the experiment. Five participants were excluded because they had previously interacted with the robot, two because they suspected the robot to be remotely controlled, and one because of eye-tracking failures occurring in all three sessions. The remaining 52 participants ($M = 38$; $F = 13$, 1 undisclosed) had an age comprised between 19 and 50 years ($M = 24.50$, $SD = 4.65$). Of them, 47 had valid gaze data for session 1 (Human: $N = 14$, Mechanical: $N = 16$, Morph: $N = 17$), 46 for session 2 (Human: $N = 15$, Mechanical: $N = 16$, Morph: $N = 15$), and 41 for session 3 (Human: $N = 17$, Mechanical: $N = 11$, Morph: $N = 13$). The study was approved by the regional ethics board, and participants were compensated with course credits for their time.

4.2. Scenario

Our experiment aimed to study people's gaze patterns in a face-to-face interaction and a joint task. We thus designed a scenario consisting of two distinct parts: a geography-themed collaborative Rapid Dialog Game (RDG) and a social chat. In the collaborative RDG-Map game (see RDG-Map game video demonstration¹), the human and the robot were tasked with identifying as many countries as possible on the world map (Paetzel and Manuvinakurike, 2019). Participants had the role of the *tutor* in this scenario. They saw a map with one country highlighted as the target. Their goal was to verbally describe this country to the robot, which acted as a *learner* with limited initial knowledge about the world map. Once the robot gained sufficient confidence about the described country, it made a guess about it and showed it on a shared screen placed in between the human and the robot. For each country correctly identified, the team received 2 points if the robot could guess the country at the first try and 1 point if it was able to guess it only at the second try. The more countries the human-robot team could identify in a given time of 10 min, the higher their score would be, and the larger the robot's knowledge base would become. The game score and the time left to score points were displayed on the shared screen positioned between the two players (for more details on the map game dynamics, consult the **Supplementary Material**).

Before and after the game, the robot engaged the human in a 2-min social chat. The chat's content varied between sessions but not between participants, and involved topics

¹<https://www.youtube.com/watch?v=U9TGrs01Am4>

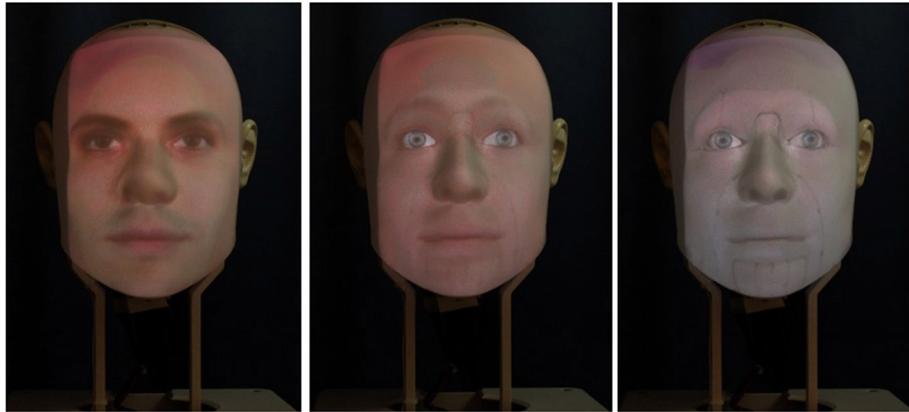


FIGURE 1 | The Furhat robot with the humanlike (**left**), mechanical (**right**), and morph (**center**) facial texture applied.

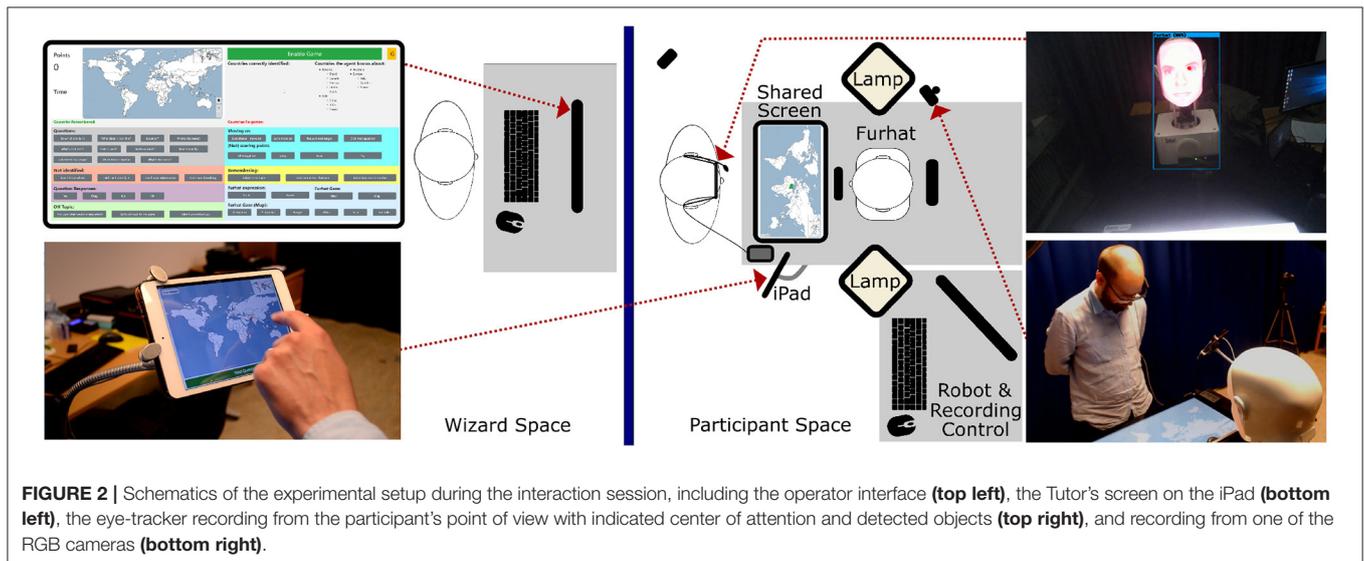


FIGURE 2 | Schematics of the experimental setup during the interaction session, including the operator interface (**top left**), the Tutor's screen on the iPad (**bottom left**), the eye-tracker recording from the participant's point of view with indicated center of attention and detected objects (**top right**), and recording from one of the RGB cameras (**bottom right**).

such as favorite games, countries that the human and the robot had visited, and future travel plans. In the second and third sessions, the robot remembered a few countries from the previous game interactions and facts from previous social chats.

4.3. Robot Embodiment and Behavior

To alter the anthropomorphic appearance of the robot while limiting confounding factors in the embodiment, we used a Furhat V1 blended robot platform (Al Moubayed et al., 2012). Furhat is a head-only robot with a semi-translucent mask on which a virtual face is projected from within. Animating the virtual face texture allows the robot to move its mouth in sync with speech, perform facial expressions, and change gaze direction. In addition, the robot's two high-torque Dynamixel servos can be used to change the head's pitch and yaw. The robot head follows the standard motion dynamics provided by

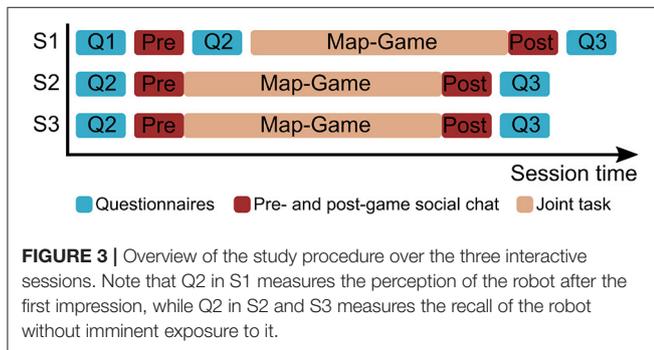
the IrisTK framework² when its pitch and yaw is altered. Taken together, the virtual animations of the face and the physical head manipulations allow to accurately direct the robot's focus of attention so it can be detected by a human interaction partner (Al Moubayed and Skantze, 2012).

To alter the perception of humanlikeness and the associated feeling of likability, we used morphing, a common approach in the literature on uncanny feelings toward artificial agents (e.g., Hanson, 2006; MacDorman et al., 2009; McDonnell et al., 2012). Three different facial textures with varying degrees of anthropomorphic features were used in our experiment (cf. **Figure 1**). The *humanlike* texture was based on the photograph of a human face. Similarly, the *mechanical* texture utilized a picture of a mechanical robot's face with parts such as screws visible in the texture. The *morph* texture was created

²<http://www.speech.kth.se/iristk/>

TABLE 1 | The distribution of Furhat's gaze between the human dialog partner, the shared screen and elsewhere, divided by interaction session, and the task-based (game) and social chat (pre- and post-game dialogue).

	Session 1			Session 2			Session 3		
	Social chat		Task	Social chat		Task	Social chat		Task
	Pre (%)	Post (%)	Game (%)	Pre (%)	Post (%)	Game (%)	Pre (%)	Post (%)	Game (%)
User	99.4	99.5	2.1	99.5	99.6	2.0	99.9	99.9	1.9
Shared screen	0.0	0.0	97.9	0.0	0.0	98.0	0.0	0.0	98.1
Somewhere else	0.6	0.5	0.0	0.5	0.4	0.0	0.1	0.1	0.0



by blending the humanlike and the mechanical face, keeping features from both of them. The particular set of facial textures utilized in this study is based on a interactive study we ran with the Furhat robot where we found that the morph robot elicited significantly higher discomfort than both the humanlike and the mechanical texture (Paetzel and Castellano, 2019). In previous work, we additionally validated the blending technique on another set of humanlike and mechanical textures and found some of the corresponding morphs to elicit significantly higher feelings of discomfort in participants compared to the original humanlike and mechanical textures (Paetzel et al., 2018).

The robot's verbal and non-verbal behavior in the interaction sessions was remote-controlled by a researcher, who followed detailed instructions to select the robot's verbal responses from a set of utterances provided by an interface (cf **Figure 2** top left, for more details, consult the **Supplementary Material**). The researcher was trained during 50 online sessions to ensure that the behavior of the robot was comparable between participants. The gaze behavior of the robot differed between the social chat and the collaborative game (cf. **Table 1**). In the social chat, the robot autonomously tracked the participant's head and kept eye contact. In the game, instead, the robot focused its gaze on the shared screen. To ensure that the behavior of the robot was perceived as natural as possible, the human controller occasionally directed the gaze of the robot to the bottom left or right to simulate thinking during the social chat. Similarly, in the joint task, the human controller directed the gaze of the robot toward the human game partner in case long periods of silence occurred. This means that, in the game context, the shared screen acted as an object of

shared attention and the iPad as an object of *exclusive attention* for the participants (cf. **Figure 2**). Moreover, it also entails that, while in the social chat participants' gaze toward the robot could be considered mutual (i.e., when the participants looked at the robot, they made eye-contact with it), in the joint task it cannot, as the robot only rarely gazed at the participants (cf. **Table 1**).

4.4. Questionnaires and Recordings

To measure participants' perception of the robot and their engagement with it and the game, we asked them to complete a series of questionnaires. Before their first interaction with the robot they filled out a demographic questionnaire (Q1). The second questionnaire (Q2) was used to capture people's perception of the robot. It contained questions about the robot's perceived *anthropomorphism* (5 items on a 5-point Likert scale from the Godspeed questionnaire, $\alpha = 0.91$; Bartneck et al., 2009), *likability* and *threat* (5-point Likert scale, likability: $\alpha = 0.83$, perceived threat $\alpha = 0.89$; Rosenthal-von der Pütten and Krämer, 2014), as well as its perceived *warmth*, *competence*, and *discomfort* (Robotic Social Attributes Scale; 18 items on a 7 point Likert scale; warmth: $\alpha = 0.92$; competence: $\alpha = 0.95$; discomfort: $\alpha = 0.90$; Carpinella et al., 2017). In the first session, Q2 was filled out immediately after the social chat with the robot to collect people's first impression of it. In the second and third session, it was instead completed before the first social chat to understand participants' recall of the robot's perception before seeing it again (cf. **Figure 3**). The final questionnaire (Q3) was filled out after the post-game social chat. It contained the same questions of Q2, but also additional scales to measure participant's *involvement* with the robot and with the game (User Engagement Questionnaire; 9 items on a 5 point-Likert scale: involvement: $\alpha = 0.71$; O'Brien and Toms, 2010).

Participants were equipped with Tobii Glasses 2 (cf. **Figure 4**), which recorded the experimental session from a first-person view with a full HD wide-angle camera. These glasses also tracked the participants' gaze direction with a sampling rate of 100 Hz. Further processing of gaze data is described in section 5. Following the Ethographic and Laban-Inspired Coding System of Engagement (ELICSE) proposed by Perugia et al. (2017, 2018), we focused on three foci of attention in the interaction: the robot, the shared screen, and the tablet, and measured the percentage of time participants gazed at each attentional focus during the different phases of the interaction session (i.e., social chats and collaborative game). To ensure the eye-tracker would not disturb

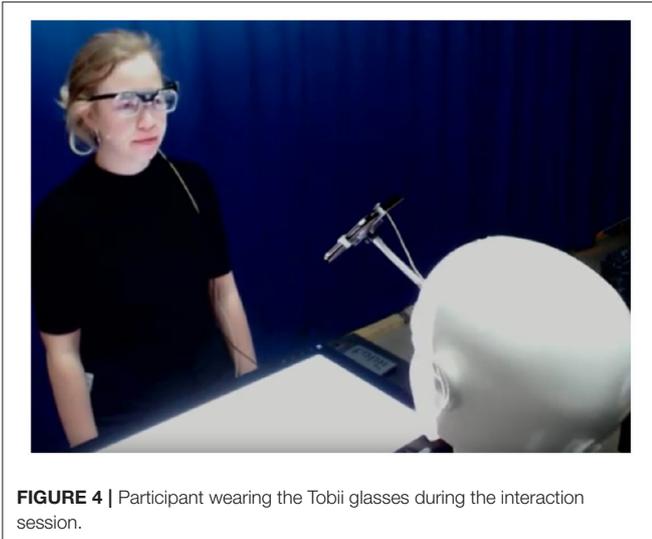


FIGURE 4 | Participant wearing the Tobii glasses during the interaction session.

participants in their interactions, we ran a pilot study with six participants (three with and three without the eye-tracking glasses). Neither the participants wearing the eye-tracker nor the control group found the recording setup intrusive. Participants' interaction with Furhat was further recorded using a close-range Sennheiser microphone, two webcams, a Kinect, and a RealSense camera. These recordings were used to answer different research questions and are hence not discussed in this paper.

4.5. Experiment Setup and Procedure

The interaction space was set up with a table on which the shared touch screen, the Furhat robot, and the iPad were placed (cf. **Figure 2**). Participants stood on one side of the table. The robot was placed in front of them roughly at the height of their eyes. The shared screen was positioned between the participant and the robot. A professional lighting system ensured even illumination for the video recordings and visibility for the robot's face texture.

During the first session (S1), participants were explained the experiment and asked to give informed consent. Then, they filled out Q1 on the iPad while the robot was still covered with a blanket. The researcher leading the experiment removed the blanket from the robot's head before manually starting the interaction. After the 2-min pre-game social chat, the robot asked participants to fill out Q2 on the iPad; then it automatically continued with the map game and the post-game social chat. At the end of the session, the robot prompted the participant to respond to Q3. The second (S2) and third sessions (S3) started with the researcher asking participants to fill in Q2 based on their memory from the previous session, before uncovering the robot. The pre-game social chat, the game interaction, and the post-game social chat were then performed without a break in between. Hence, while Q3 was always filled out at the same time, immediately after the post-game social chat, Q2 was completed *after* the pre-game social chat in S1, and *before* it in S2 and S3 (cf. **Figure 3**). While participants responded to a questionnaire,

the robot displayed idling behavior that involved looking around in the room and away from the human interaction partner. Participants were fully debriefed about the purpose of the study after the entire experiment was completed.

5. DATA PROCESSING

To understand what object participants were focusing on, we developed an object detector for the first-person video stream from the wearable eye-tracker. The implementation of the object detector was based on the open-source neural network framework Darknet, which uses the real-time object detector YOLOv4 (Bochkovskiy et al., 2020). For the purpose of this study, we used a version of YOLOv4 pre-trained on the MS COCO data set consisting of objects such as cars, tv screens, and people (Lin et al., 2014), and added 409 labeled images of tablets and the Furhat robot. The resulting model achieved a mAP of 89.06%, with the shared touchscreen, robot, and tablet having an AP of 85.07, 89.35, and 92.75%, respectively.

To run the analyses on gaze, we extracted the *percentage of gaze* directed to the robot, screen, and tablet from each interaction phase. We then compared every frame of the gaze coordinates provided by the Tobii eye-tracking system with the objects detected in the video stream and labeled them as either inside of the bounding box of the robot, screen or tablet, or "somewhere else." The Tobii system failed to detect participants' pupils on average on 11.55% of the frames ($SD = 8.9\%$), in which case the frame was annotated as "Not applicable." Interaction phases containing more than 50% of undetected frames were excluded from the analysis. To correct for inaccuracies due to the inexact positioning of the bounding boxes in the first person video, we applied a filter to the resulting object annotations. The filtering algorithm detected one or two consecutive frames labeled as outside the bounding box of an object occurring in the middle of a larger block of frames detected as inside the bounding box of that object. If the distance between the frames labeled as outside and those labeled as inside the bounding box was lower or equal to 110.14 pixels (5% of the max. video distance), we changed the original label of the outlier frames to the label of the surrounding block of frames.

Two annotators manually labeled three of the videos frame-by-frame using the software ELAN 5.9. The inter-rater agreement between the two annotators, which was calculated on one video, was excellent ($\kappa = 0.98$; Holle and Rein, 2015). When comparing the automated annotations to the manual ones, the system achieved a similarly excellent average κ of 0.97.

6. RESULTS

In the following, we use: (i) *perception of the robot* to refer to the subscales anthropomorphism, perceived threat, likability, warmth, competence, and discomfort; (ii) *engagement* to refer to the subscales involvement with the game and involvement with the robot; and (iii) *task performance* to refer to participants' game score. Moreover, when it comes to gaze metrics, we use:

(a) *mutual gaze* to refer to the percentage of gaze directed to the robot during the pre- and post-game social chats; and (b) *gaze patterns in the joint task* to refer to the percentage of gaze toward the robot, screen, and tablet during the game interaction. All dependent variables used for the statistical analyses were normally distributed and met the equality of variance assumption. The subscales related to participants' perception of the robot and their involvement with the game and the robot were always used in their original form.

To correct for multiple testing in the univariate tests and between-subjects effects following up a MANOVA, we opted for a Holms-Bonferroni correction (Holm, 1979). Holm's method enabled us to control family-wise error rates (FWER) while at the same time keeping an optimal statistical power (Haynes, 2013). To adjust the *p*-values of *post-hoc* pairwise comparisons following the univariate and between-subjects effects, instead, we used a Bonferroni correction (Bonferroni, 1936). This more conservative approach was meant to compensate for the further iteration of analysis. In the classical Bonferroni test, the alpha levels obtained from the statistical analyses are compared to the one resulting from the following correction: $\frac{\alpha}{n}$, where *n* is the number of tests performed, and α is usually 0.05 (Bonferroni, 1936). In the Holm-Bonferroni correction, or sequentially rejective Bonferroni test, instead, the obtained alpha levels are first ranked and then sequentially compared to the values resulting from the following equations (Holm, 1979):

$$\frac{\alpha}{n}, \frac{\alpha}{n-1}, \frac{\alpha}{n-2}, \dots, \frac{\alpha}{1} \quad (1)$$

6.1. Manipulation Check

In order to check whether we succeeded in manipulating the robot's humanlikeness, we carried out a repeated measures MANOVA with humanlikeness as between-subject factor (humanlike, mechanical, and morph), interaction session as within-subject factor (S1, S2, and S3), and perception of the robot (Q3) as dependent variable (i.e., anthropomorphism, likability, warmth, competence, threat, and discomfort). The results did not disclose a significant main effect of humanlikeness on the linear composite of the six dependent variables [$F_{(12, 72)} = 1.475$, $p = 0.154$, $\eta p^2 = 0.197$] nor a significant interaction effect of humanlikeness and interaction session [$F_{(24, 60)} = 0.678$, $p = 0.853$, $\eta p^2 = 0.213$]. However, they showed a significant main effect of interaction session on the composite of the dependent variables [$F_{(12, 29)} = 4.726$, $p < 0.001$, $\eta p^2 = 0.662$].

Considering a Holm-Bonferroni correction (cf. Table 3), the univariate analyses disclosed a main effect of interaction session on perceived threat [$F_{(2, 80)} = 4.984$, $p = 0.009$, $\eta p^2 = 0.111$], and discomfort [$F_{(2, 80)} = 12.920$, $p < 0.001$, $\eta p^2 = 0.244$], but not on anthropomorphism [$F_{(2, 80)} = 0.179$, $p = 0.837$, $\eta p^2 = 0.004$], likability [$F_{(2, 80)} = 3.897$, $p = 0.024$, $\eta p^2 = 0.089$], warmth [$F_{(2, 80)} = 1.533$, $p = 0.222$, $\eta p^2 = 0.037$], and competence [$F_{(2, 80)} = 2.968$, $p = 0.057$, $\eta p^2 = 0.069$]. In line with previous work (Paetzl et al., 2020), *Post-hoc* analyses with a Bonferroni correction showed that perceived threat and discomfort did not stabilize over time. Perceived threat decreased between S1 and S3 ($p = 0.038$; S1–S2: $p = 0.377$; S2–S3:

TABLE 2 | Mean (*M*) and standard deviation (*SD*) of the different perceptual dimensions (Q3) per session.

	Manipulation check: descriptive statistics (session)					
	Session 1		Session 2		Session 3	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Anthropomorphism	3.335	0.671	3.330	0.818	3.297	0.865
Likability	3.163	0.770	3.354	0.821	3.316	0.786
Warmth	4.132	1.165	4.287	1.314	4.140	1.438
Competence	4.919	1.079	5.128	1.041	4.861	1.263
Threat	1.870	0.638	1.772	0.555	1.6744	0.529
Discomfort	2.019	0.705	1.771	0.595	1.643	0.573

$p = 0.101$), and discomfort between S1 and S2 ($p = 0.006$) and between S1 and S3 ($p = 0.001$, S2–S3: $p = 0.054$; cf. Table 2 for the descriptive statistics).

The main effect of humanlikeness on the linear composite of the six dependent variables (i.e., anthropomorphism, likability, warmth, competence, perceived threat, and discomfort) was not significant. Nevertheless, we proceeded to check the between-subjects effects. This was because the perceptual dimensions we included in the MANOVA were in a complex relationship with each other (anthropomorphism, likability, warmth, competence are positively correlated with each other but negatively correlated with perceived threat and discomfort), hence performing analyses exclusively focusing on their linear composite could have disguised crucial underlying effects. Between-subject effects with a Holm-Bonferroni correction (cf. Table 3) highlighted a significant main effect of humanlikeness on anthropomorphism [$F_{(2,40)} = 6.986$, $p = 0.002$, $\eta p^2 = 0.259$], likability [$F_{(2,40)} = 5.441$, $p = 0.008$, $\eta p^2 = 0.214$], warmth [$F_{(2,40)} = 8.550$, $p = 0.001$, $\eta p^2 = 0.299$], and competence [$F_{(2,40)} = 8.694$, $p = 0.001$, $\eta p^2 = 0.303$], but not on perceived threat [$F_{(2,40)} = 0.003$, $p = 0.997$, $\eta p^2 < 0.001$] and discomfort [$F_{(2,40)} = 0.013$, $p = 0.987$, $\eta p^2 = 0.001$]. *Post-hoc* pairwise comparisons with a Bonferroni correction revealed that no difference in perception was present between the morph and the mechanical robot (anthropomorphism: $p = 1.00$; warmth: $p = 0.942$; likability: $p = 0.833$; competence: $p = 1.00$; discomfort: $p = 1.00$; perceived threat: $p = 1.00$). However, they disclosed that the humanlike robot was perceived as significantly more anthropomorphic ($p = 0.004$), warm ($p = 0.001$), likable ($p = 0.008$), and competent ($p = 0.001$) than the morph (cf. Table 4 for the descriptive statistics) and significantly more anthropomorphic ($p = 0.020$), warm ($p = 0.021$), and competent ($p = 0.015$) than the mechanical robot (likability: $p = 0.142$, cf. Table 4 for the descriptive statistics).

6.1.1. Discussion of Manipulation Check

In summary, perceived threat and discomfort, which should have varied due to changes in the robot's level of humanlikeness and the presence of mismatching cues in the morph robot (see Kätsyri et al., 2015), did not change as expected. Nevertheless, the results of the manipulation check show that the robot's humanlikeness

TABLE 3 | Ranked p -values for the univariate analysis and between-subjects effects referring to the manipulation check with corresponding cut-off p -values due to Holm-Bonferroni correction.

Manipulation check: Holm-Bonferroni correction					
Univariate analyses			Between-subjects effects		
Ranked scale	p -value	cut-off p	Ranked scale	p -value	cut-off p
Discomfort	*<0.001	0.008	Competence	*0.001	0.008
Threat	*0.009	0.010	Warmth	*0.001	0.010
<i>Likability</i>	<i>*0.024</i>	<i>0.013</i>	Anthropomorphism	*0.002	0.013
Competence	0.057	0.017	Likability	*0.008	0.017
Warmth	0.222	0.025	Discomfort	0.987	0.025
Anthropomorphism	0.887	0.050	Threat	0.997	0.050

The asterisk and bold type indicate the significant p -values after correction. The asterisk and italic type indicate the significant p -values before correction.

TABLE 4 | Mean (M) and standard deviation (SD) of the different perceptual dimensions (Q3) per level of humanlikeness.

Manipulation check: descriptive statistics (humanlikeness)						
	Humanlike		Morph		Mechanical	
	M	SD	M	SD	M	SD
Anthropomorphism	3.761	0.631	2.969	0.631	3.097	0.631
Likability	3.675	0.676	2.872	0.677	3.164	0.677
Warmth	4.993	1.068	3.444	1.067	3.872	1.067
Competence	5.654	0.891	4.368	0.891	4.675	0.891
Threat	1.773	0.540	1.764	0.541	1.779	0.541
Discomfort	1.824	0.577	1.791	0.577	1.816	0.577

significantly varied across conditions. This is indicated by the significant differences in anthropomorphism between the morph and the humanlike robot, and between the mechanical robot and the humanlike one, but also by changes in perceptual dimensions known to be related to a robot's humanlikeness, such as likability, warmth, and competence. The lack of proper differentiation between the morph and the mechanical robot in terms of anthropomorphism can be ascribed to the many facial features the two robots had in common (cf. **Figure 1**). In the future, the humanlike characteristics of the morph robot should be strengthened to increase its recognizability and enhance its ambiguity and hence its uncanniness. In conclusion, given that the core independent variable of our study, the humanlikeness of the robot, varied as expected, we did not consider that the lack of significant differences in perceived threat and discomfort could undermine the relevance of further analyses. Indeed, we consider these two dimensions to be of further relevance to our analysis because they were the only perceptual dimensions to significantly change over time.

6.2. Preliminary Analyses: Engagement and Task Performance

To understand the effects of our study design on engagement and task performance, we conducted a repeated measure MANOVA with the same independent variables (humanlikeness as between-subjects factor; interaction session as within-subject factor) and involvement with the robot, involvement with the

game, and task performance (i.e., score at the game) as dependent variables. Results disclosed a significant main effect of interaction session [$F_{(6, 35)} = 9.005, p < 0.001, \eta p^2 = 0.607$] and a trend main effect of humanlikeness [$F_{(6, 78)} = 1.936, p = 0.085, \eta p^2 = 0.130$] on the linear composite of the three dependent variables. No interaction effect between humanlikeness and interaction session was present [$F_{(12, 72)} = 0.866, p = 0.584, \eta p^2 = 0.126$].

Considering a Holm-Bonferroni correction (cf. **Table 5**), the univariate analyses showed a significant main effect of interaction session on task performance [$F_{(2, 80)} = 37.208, p < 0.001, \eta p^2 = 0.482$] but not on involvement with the robot [$F_{(2, 80)} = 0.576, p = 0.564, \eta p^2 = 0.014$] and the game [$F_{(2, 80)} = 1.606, p = 0.207, \eta p^2 = 0.039$]. *Post-hoc* analyses with p -values adjusted with a Bonferroni correction disclosed a significant difference in task performance between S1 ($M = 22.720, SD = 9.59$) and S2 ($M = 28.419, SD = 10.15, p < 0.001$), S2 and S3 ($M = 30.558, SD = 10.24, p = 0.019$), and S1 and S3 ($p < 0.001$). Interestingly, albeit not significant, engagement was higher in the humanlike robot condition compared to the morph condition, while task performance was higher for the morph robot with respect to the humanlike robot (cf. **Figure 5**).

6.3. Analyses of the Research Questions

6.3.1. Mutual Gaze as Predictor of Perceptions of Robots (RQ1)

We ran a number of regression analyses using mutual gaze in the post-game social chat as independent variable and the dimensions of perception of the robot (Q3) as dependent variables. Mutual gaze during the social chat was not a significant predictor of anthropomorphism and competence (cf. **Table 6**). However, it was a significant negative predictor of perceived threat and discomfort, and a significant positive predictor of likability and warmth (cf. **Table 6**). Hence, we can conclude that *the less people gazed at the robot during the social chat, the less they liked the robot, and the more they perceived it as uncanny*.

6.3.2. Gaze Patterns as Predictors of Engagement and Task Performance (RQ2)

To understand whether participants' gaze patterns during the joint task were predictors of their engagement and task performance, we ran separate regression analyses using the

TABLE 5 | Ranked *p*-values for the univariate analysis and between-subjects effects referring to the preliminary analyses with corresponding cut-off *p*-values due to Holm-Bonferroni correction.

Preliminary analyses: Holm-Bonferroni correction					
Univariate analyses			Between-subjects effects		
Ranked scale	<i>p</i> -value	Cut-off <i>p</i>	Ranked scale	<i>p</i> -value	Cut-off <i>p</i>
Task performance	*<0.001	0.017	<i>Involvement Robot</i>	<i>*0.037</i>	<i>0.017</i>
Involvement game	0.207	0.025	Involvement game	0.075	0.025
Involvement robot	0.576	0.050	Task performance	0.131	0.050

The asterisk and bold type indicate the significant *p*-values after correction. The asterisk and italic type indicate the significant *p*-values before correction.

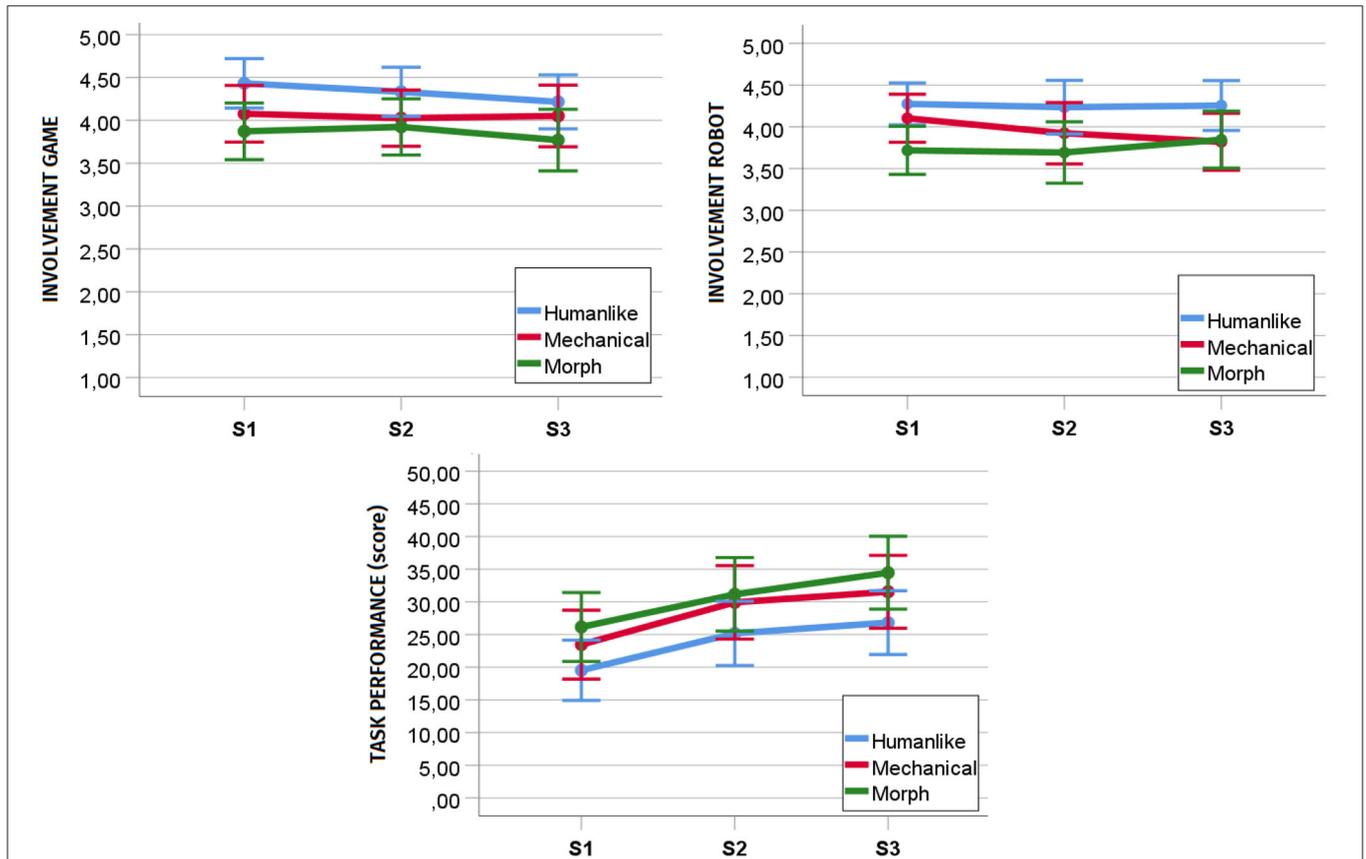


FIGURE 5 | Development of involvement with the game, involvement with the robot, and task performance (participants' score at the game) over repeated sessions. S1, Session 1; S2, Session 2; S3, Session 3.

percentage of gaze directed toward the robot, toward the screen, and toward the tablet during the game interaction as predictors and involvement with the robot and with the game, and task performance as dependent variables.

The percentage of gaze directed to the robot during the game was not a significant predictor of involvement with the robot, nor of involvement with the game (cf. Table 7). However, it was a significant negative predictor of task performance, meaning that *the more participants looked at the robot during the game, the less they scored at the game.*

The percentage of gaze directed to the screen during the game was not a significant predictor of involvement with the robot.

However, it was a significant predictor of involvement with the game and especially of task performance (cf. Table 7). This indicates that *the more participants focused on the object of shared attention (the screen), the more they were engaged with the game and the higher they scored at the game.*

Finally, the percentage of gaze directed to the tablet during the game was not a significant predictor of involvement with the robot, but it was a significant predictor of involvement with the game and task performance (cf. Table 7). As opposed to the percentage of gaze directed the screen, *the more participants looked at the object of exclusive attention (the tablet), the less they were engaged with the game and the lower they scored at the game.*

TABLE 6 | Results of the regression analyses performed for RQ1.

Regression analyses (RQ1): results				
Predictor: mutual gaze				
Dependent variable	β	$t_{(132)}$	p -value	r^2
Anthropomorphism	0.064	0.733	0.465	0.004
Likability	0.219	2.574	0.011	0.048
Warmth	0.172	2.000	0.048	0.030
Competence	0.015	0.177	0.860	<0.001
Threat	-0.216	-2.534	0.012	0.047
Discomfort	-0.221	-2.592	0.011	0.049

In bold, the significant regressions.

TABLE 7 | Results of the regression analyses performed for RQ2.

Regression analyses (RQ2): results				
Predictor: gaze to robot				
Dependent variable	β	$t_{(130)}$	p -value	r^2
Involvement robot	0.021	0.243	0.809	<0.001
Involvement game	-0.041	-0.470	0.639	0.002
Task performance	-0.249	-2.917	0.004	0.062
Predictor: gaze to screen				
Involvement robot	0.128	1.466	0.146	0.016
Involvement game	0.192	2.220	0.028	0.037
Task performance	0.305	3.643	<0.001	0.093
Predictor: gaze to tablet				
Involvement robot	-0.118	-1.351	0.179	0.014
Involvement game	-0.172	-1.983	0.049	0.030
Task performance	-0.248	-2.907	0.004	0.061

In bold, the significant regressions.

6.3.3. Effect of Exposure, Interaction Session, and Humanlikeness on Mutual Gaze (RQ3, RQ5a)

To understand how mutual gaze developed over time within and between interactions, we performed a repeated measures ANOVA with the robot's humanlikeness as between-subject factor (humanlike, mechanical, and morph), interaction session (S1, S2, and S3) and game exposure (pre- and post-game) as within-subject factors, and mutual gaze as dependent variable. We found a significant main effect of game exposure [$F_{(1, 30)} = 23.515, p < 0.001, \eta p^2 = 0.439$] and an interaction effect of game exposure and interaction session on mutual gaze in the social chats [$F_{(2, 29)} = 20.999, p < 0.001, \eta p^2 = 0.592$]. However, we did not find any significant main effect of interaction session [$F_{(2, 29)} = 2.107, p = 0.140, \eta p^2 = 0.127$] and humanlikeness on mutual gaze [$F_{(2, 30)} = 0.320, p = 0.728, \eta p^2 = 0.021$], nor any interaction effect of interaction session and humanlikeness [$F_{(4, 60)} = 0.806, p = 0.526, \eta p^2 = 0.051$], game exposure and humanlikeness [$F_{(2, 30)} = 0.146, p = 0.864, \eta p^2 = 0.010$], and interaction session, game exposure and humanlikeness [$F_{(4, 60)} = 0.958, p = 0.437, \eta p^2 = 0.060$].

Post-hoc analyses with a Bonferroni correction showed a general decrease in mutual gaze from the pre- ($M = 0.714, SD =$

0.126) to the post-game social chat ($M = 0.630, SD = 0.144, p < 0.001$). Follow-up separate univariate analyses disclosed that the mutual gaze toward the robot significantly decreased [$F_{(1, 44)} = 48.985, p < 0.001, \eta p^2 = 0.527$] from the pre- ($M = 0.729, SD = 0.123$) to the post-game social chat in S1 ($M = 0.575, SD = 0.157$) and in S2 [$F_{(1, 42)} = 25.398, p < 0.001, \eta p^2 = 0.377$; pre: $M = 0.723, SD = 0.126$; post: $M = 0.638, SD = 0.147$], but not in S3 [$F_{(1, 38)} = 0.002, p = 0.968, \eta p^2 = 0.00$; pre: $M = 0.688, SD = 0.172$; post: $M = 0.679, SD = 0.172$, cf. **Figure 6**]. Interestingly, in the last session, the amount of gaze toward the robot was almost same in the pre and post-game social chats (cf. **Figure 6**). This might indicate that *the more participants interacted with the robot, the more they habituated to it and the more their gaze patterns stabilized*.

6.3.4. Effect of Interaction Session and Humanlikeness on Gaze Patterns (RQ4, RQ5b)

To understand whether participants' gaze patterns in the joint task changed across repeated interactions, we performed a repeated measures MANOVA with humanlikeness as between-subject factor (humanlike, mechanical, and morph), interaction session as within-subject factor (S1, S2, and S3), and gaze patterns as dependent variables (i.e., percentage of gaze directed to the robot, the screen, and the tablet). We did not find a significant main effect of humanlikeness [$F_{(6, 58)} = 0.422, p = 0.861, \eta p^2 = 0.042$] nor an interaction effect of humanlikeness and interaction session on the dependent variables [$F_{(12, 52)} = 0.961, p = 0.496, \eta p^2 = 0.182$]. However, the results disclosed a significant main effect of interaction session on the dependent variables [$F_{(6, 25)} = 2.543, p = 0.046, \eta p^2 = 0.379$].

Considering a Holm-Bonferroni correction (cf. **Table 8**), further univariate tests showed a significant main effect of time on the percentage of gaze directed to the screen [$F_{(2, 60)} = 7.143, p = 0.002, \eta p^2 = 1, 928$] and the tablet [$F_{(2, 60)} = 9.666, p < 0.001, \eta p^2 = 0.244$], but not of the percentage of gaze directed to the robot [$F_{(2, 60)} = 2.202, p = 0.119, \eta p^2 = 0.068$]. *Post-hoc* analyses with *p*-values adjusted with a Bonferroni correction revealed a significant decrease in the percentage of gaze directed to the screen in the joint task from S1 ($M = 0.423, SD = 0.185$) to S3 ($M = 0.371, SD = 0.191, p = 0.023$), and from S2 ($M = 0.425, SD = .190$) to S3 ($p = .007$), but not from S1 to S2 ($p = 1.00$, cf. **Figure 7**). Moreover, they revealed a significant increase in the percentage of gaze directed to the tablet in the joint task from S1 ($M = 0.374, SD = 0.172$) to S3 ($M = 0.458, SD = 0.183, p = 0.004$) and from S2 ($M = 0.390, SD = 0.163$) to S3 ($p = 0.002$), but not from S1 to S2 ($p = 1.00$, cf. **Figure 7**). This seems to suggest that *the gaze patterns in the joint task changed over time with a decrease in gaze toward the object of shared attention making space for an increase in gaze toward the object of exclusive attention*.

7. DISCUSSION

7.1. Mutual Gaze and Uncanniness (RQ1)

In our experiment, mutual gaze in a social chat was a negative predictor of uncanniness (i.e., perceived threat and discomfort) and a positive predictor of likability (i.e., likability and warmth).

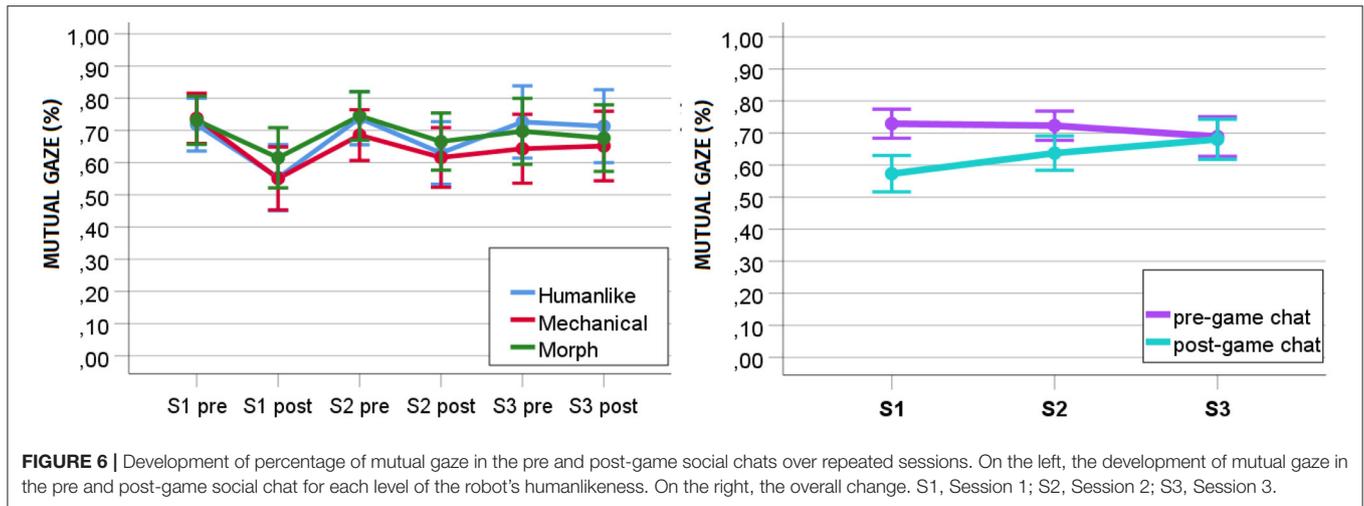


TABLE 8 | Ranked p -values for the follow-up univariate analysis performed for RQ4 with corresponding cut-off p -values due to Holm-Bonferroni correction.

RQ4: Holm-Bonferroni correction		
Ranked scale	p -value	cut-off p
Gaze table	*<0.001	0.017
Gaze screen	*0.002	0.025
Gaze robot	0.119	0.050

In bold, the significant p -values after correction. The asterisk indicates the significant p -values before correction.

The negative relation between mutual gaze and uncanniness and the positive relation between mutual gaze and likability lend support to the mutual gaze-liking hypothesis. Indeed, *robots perceived as uncanny seem to elicit gaze aversion, whereas robots perceived as likable attract higher gaze allocation*. In this sense, our work extends previous findings on the relationship between gaze and uncanniness to the context of face-to-face interactions with robots. Moreover, it *suggests that people's mutual gaze in an interaction with robots can be used as an implicit and continuous measure of a robot's uncanniness and likability*. Future work should corroborate this finding in a less exploratory way, by exposing people to robots explicitly manipulated in their level of uncanniness and likability and assessing whether our results still hold.

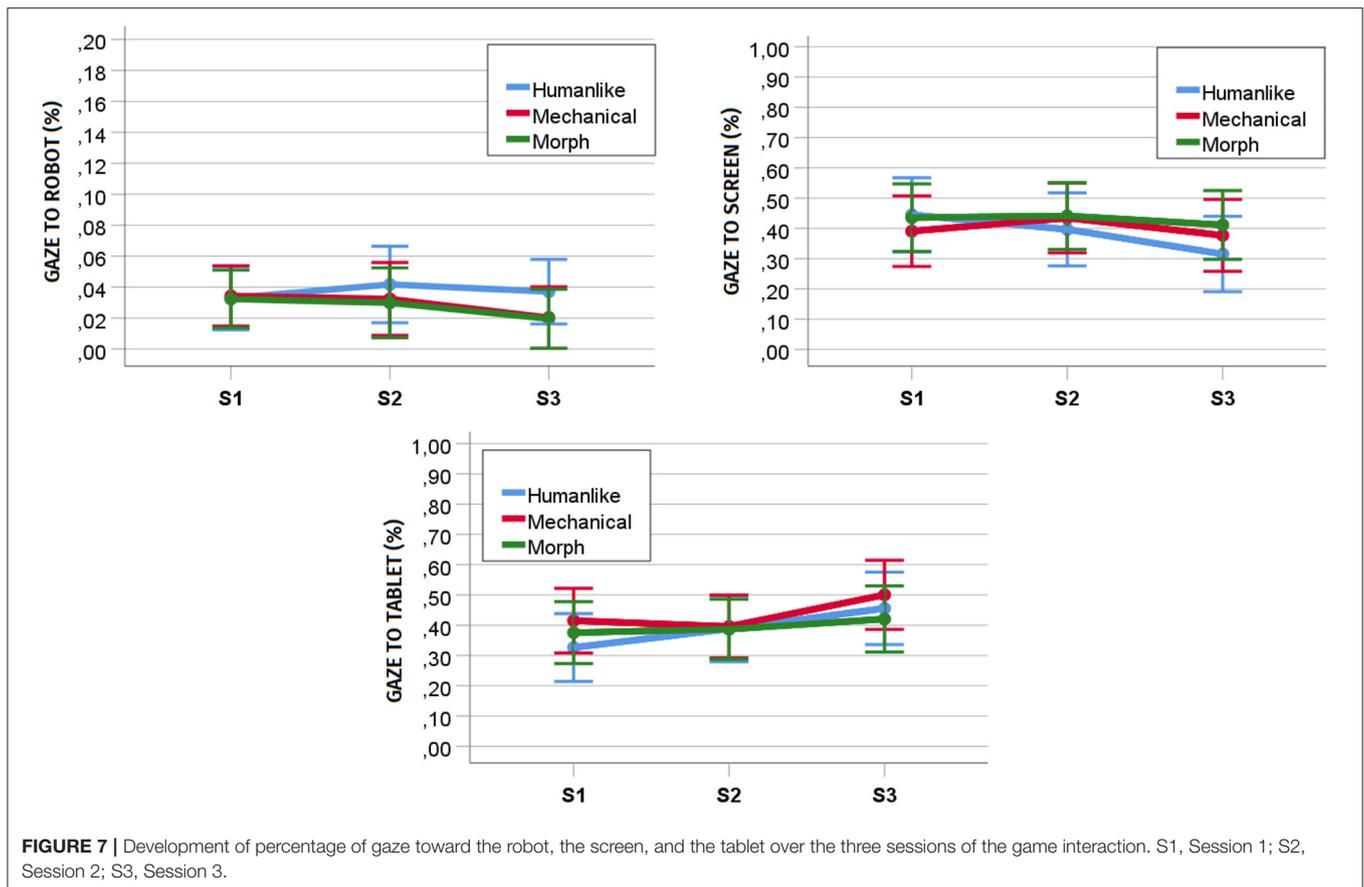
7.2. Shared Gaze, Engagement and Task Performance (RQ2)

Participants that gazed at the screen longer during the game interaction felt a higher involvement with the game and performed better. As the screen acted as the object of shared attention in this study, these results entail that *the more participants shared the focus of their attention with the robot, the more they felt involved with the game, and the better they performed*. This claim is further supported by the fact that

the gaze directed toward the object of exclusive attention (i.e., the tablet) negatively predicted involvement with the game and task performance, and the gaze directed to the robot negatively predicted participants' performance. Overall, we can state that gaze patterns in a joint task predict task performance and involvement with the game and that *in a joint task involving tangible artifacts (e.g., the screen), shared attention signals higher involvement with the task and can predict a better performance*. On the contrary, *gaze directed to the robot and the object of exclusive attention (e.g., the tablet) are markers of disengagement with the task and poorer task performance*.

In contrast with most HRI literature that employed gaze toward the robot as one of the core metrics of social engagement in a joint task, we did not find a relationship between gaze toward the robot and participants' perceived involvement with it. This confirms our suspect that in a joint task, the gaze allocated to the robot is not a precise measure of people's syntony with it because it is hindered by participants' willingness to complete the task. Combining this result with our findings on mutual gaze, we posit that *gaze toward the robot does indicate social engagement, but only in interactions that do not involve the use of tangible artifacts, for instance, in face-to-face social dialogs*. Joint tasks involving tangible artifacts call for the allocation of attentional resources to the object where the activity takes place (in our case, the shared screen) rather than to the agent with which the activity is performed. Hence, *we argue that, in these tasks, one can feel involved with the robot at a subjective-experiential level even without overtly expressing this involvement at a behavioral level*. Future work should focus on testing these preliminary findings in further joint tasks and see if they still hold.

In this study, we found that: (1) the amount of mutual gaze in the social chat was a negative predictor of uncanniness and a positive predictor of likability, (2) the gaze directed to the robot in the joint task was a negative predictor of task performance, and (3) the percentage of gaze directed to the screen was a positive predictor of task performance. Altogether, *this seems to suggest that robots that are perceived as less likable might be more suitable*



for joint tasks, as they attract less attention and hence help the player stay focused on the activity. It would be interesting to understand whether we could leverage on a robot's likability to find a trade-off between engagement and task performance in joint activities.

7.3. The Development of Mutual Gaze Over Time (RQ3)

We found participants' mutual gaze in the face-to-face conversation with the robot to change over time. It decreased between the pre- and post-game social chat in sessions 1 and 2, but not in session 3. The descriptive statistics reveal that mutual gaze in the third pre- and post-game chats stabilizes close to the values of the pre-game conversations of session 1 and 2. In line with the questionnaire results on the perception of robots, which revealed that perceived threat and discomfort were the last perceptual dimensions to stabilize over time, in this study, we found that mutual gaze, a negative predictor of perceived threat and discomfort, stabilized only at the third interaction session. Consistent with self-reports from participants, which showed that uncanniness reduced over time, we found mutual gaze, a negative predictor of uncanniness, to increase across sessions. These results seem to suggest that *mutual gaze can be used to monitor the development of uncanny feelings toward a robot over time.*

The decrease in mutual gaze between the pre- and post-game social chats of sessions 1 and 2 might be related to the robot's novelty. Indeed, it seems to suggest that participants look more at a robot when meeting it for the first time and after a period of zero exposure and that they gaze progressively less at the robot the more they become familiar with it. However, the reduction in mutual gaze *within* the interaction session was accompanied by an increase in mutual gaze *across* interaction sessions. This makes it challenging to draw conclusions on the role of the robot's novelty on mutual gaze. Future research should specifically investigate how the robot's novelty affects the amount of mutual gaze it attracts and how the interplay of novelty and uncanniness influences mutual gaze within and between interaction sessions.

7.4. The Development of Gaze Patterns in a Joint Task Over Time (RQ4)

As opposed to previous work finding a decrease in the gaze directed toward the robot over multiple sessions of a joint task (Serholt and Barendregt, 2016), in our study, we found that the percentage of gaze directed to the robot during the map game was stable. This is in line with the results from the questionnaires (i.e., involvement with the robot). This result is positive as it shows that the map game is interesting enough to sustain participants' engagement with the robot over time. As gaze toward the robot

during a joint task seems to be a significant predictor of poor task performance, preventing an increase in the attention the robot attracts across repeated interactions is crucial to ensure that the educational game fulfills its pedagogical objectives.

In contrast with the gaze toward the robot, the percentage of gaze directed to the screen or the tablet changed over time, with the former decreasing and the latter increasing over the last two sessions. This result is interesting. Indeed, the progressive improvement in task performance across sessions, the positive relationship between gaze toward the screen and task performance, and the negative relationship between gaze toward the tablet and task performance would have suggested an inverse development of gaze patterns over time. Hence, we suppose that this change in gaze patterns might capture the slight decrease in involvement with the game shown by the questionnaires, which eventually did not reach significance. However, it might also indicate that, as participant grew more confident with the game and settled for a strategy to score points in the last sessions, they felt more comfortable in abandoning the main support tool offered by the game (i.e., shared screen). Future research should investigate more thoroughly how gaze allocation to the objects of attention included in a joint task changes over time, especially as a consequence of the progressive increase in participants' task expertise.

7.5. The Effect of Humanlikeness on Mutual Gaze in a Social Chat (RQ5a)

The three facial textures that we applied to the robot varied in terms of positive perceptions but not in uncanniness. As mutual gaze predicted only two perceptual dimensions (i.e., likability, warmth) out of the four positive ones that varied, the lack of a main effect of humanlikeness on mutual gaze does not surprise. We assume that a less subtle manipulation of humanlikeness will be more likely to influence the gaze allocation toward the robot in a social chat, and strongly advise future research to move in this direction. At the same time, we also recommend to keep the embodiment features of the robot as consistent as possible across conditions to limit the influence of other confounding factors.

7.6. The Effect of Humanlikeness on Gaze Patterns in a Joint Task (RQ5b)

The significance values in **Table 5** and the graphs in **Figure 5** show a trend difference between the morph and the humanlike robot in terms of involvement with the robot. However, we did not find a significant difference between the humanlike and the morph robot in terms of gaze allocation. This result is particularly interesting as it corroborates our hypothesis that, in joint tasks, the involvement with the robot might be felt at a subjective/experiential level rather than expressed at a behavioral level with gaze. This might be especially true for games with time constraints. Indeed, in this context, the time pressure set by the game and the pace that derives from it might leave little room for participants to focus their gaze on the robot. Future work should further investigate this line of thought by exposing participants to joint tasks differing in time constraint and investigating at which

level of time pressure the engagement with the robot ceases to be expressed behaviorally.

7.7. Limitations

While we highlight the contribution of the present exploratory study on the usage of gaze as an implicit measure of robot perception and task performance, we also acknowledge a number of limitations. For instance, the manipulation of the robot's humanlikeness in our experiment did not work as expected. Indeed, participants did not perceive the mechanical and the morph robot as differing in anthropomorphism, and Furhat's facial textures did not vary in perceived uncanniness. To overcome this drawback, we plan to add more anthropomorphic features to the morph texture in the future. Another potential limitation of the study lays in the remote-controlled nature of the robot's interactive capabilities. While participants were not aware of the robot being controlled by a human until they were fully debriefed, this might have set wrong (i.e., unrealistically high) expectations on the robot's abilities. We are currently working on a fully autonomous version of the map game, which we plan to deploy in future studies to confirm our findings. Third, although we found a large effect size for all significant analyses, future work would benefit from a larger and more heterogeneous group of participants, both in terms of background and gender. As most of the participants in this study identified themselves as male and came from a computer science background, our results might report the perspective of a limited group of users and thus need replication. Moreover, the study we performed was set in a lab environment, a context that grants a lot of control over confounding variables. Further research should focus on replicating this study in real-life scenarios where the collection of gaze data is more complex and environmental factors, such as light conditions, might intrude first-person object-recognition and hence the automatic annotation of gaze. Finally, albeit the participants involved in the pilot did not perceive the Tobii eye-tracking glasses as intrusive, some might have felt uncomfortable wearing them. Further research should hence explore the feasibility of stationary eye-trackers in similar scenarios and compare their accuracy in detecting gaze direction.

8. CONCLUSION

In this paper, participants took part in three interaction sessions with a robot varying in humanlikeness. In each session, they played a collaborative game with the robot and engaged in a brief social chat before and after the game. We gauged their gaze direction in both types of interaction and used regression analyses to relate it with measures of perception and engagement. Results suggest that mutual gaze toward a robot in a social chat is related to perceptions of uncanniness, and the gaze directed to the robot in a joint task is a predictor of poor task performance. Moreover, they show that mutual gaze in a social chat changes across repeated interaction sessions, and so do participants' gaze patterns in a joint task. These findings are crucial for the field of HRI as they highlight that gaze can be used as an implicit measure of people's perceptions of robots in a face-to-face interaction, and of engagement and task performance in a collaborative game.

DATA AVAILABILITY STATEMENT

The dataset presented in this article is not readily available because it contains personal information on the participants. An anonymized version of the data supporting the conclusions of this article can be made available by the authors upon request.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by Regionala etikprovningarnamnden Uppsala Dnr 2018/503. The patients/participants provided their written informed consent to participate in this study. Written informed consent was obtained from the individual(s) for the publication of any potentially identifiable images or data included in this article.

AUTHOR CONTRIBUTIONS

GP wrote the paper, formulated the research questions, conceived and designed the study, collected the data, annotated the videos for the inter-rater agreement, performed the statistical analyses, and interpreted the results. MP-P significantly contributed to the writing and critical revision of the paper, conceived and designed the study, wrote the program for the robot interaction, collected

the data, processed the gaze annotations, and interpreted the results. MA developed the gaze annotation tool, annotated the videos for the inter-rater agreement, and contributed to the writing of section 5 of the paper. GC read and gave comments on the final version of the paper. All authors contributed to the article and approved the submitted version.

FUNDING

This work was supported by the Swedish Foundation for Strategic Research under the COIN project (RIT15-0133).

ACKNOWLEDGMENTS

A pre-print of this article has been made available on arXiv (Perugia et al., 2021). We would like to thank Sebastian Walkötter for **Figure 1**, Robert Kessler for assisting with the technical setup of the map game, and Ramesh Manuvinakurike for his work on the game dynamics.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frobt.2021.645956/full#supplementary-material>

REFERENCES

- Admoni, H., and Scassellati, B. (2017). Social eye gaze in human-robot interaction: a review. *J. Hum. Robot Interact.* 6, 25–63. doi: 10.5898/JHRI.6.1.Admoni
- Ahmad, M. I., and Mubin, O. (2018). "Emotion and memory model to promote mathematics learning-an exploratory long-term study," in *Proceedings of the 6th International Conference on Human-Agent Interaction* (Southampton), 214–221. doi: 10.1145/3284432.3284451
- Ahmad, M. I., Mubin, O., and Orlando, J. (2017). Adaptive social robot for sustaining social engagement during long-term children-robot interaction. *Int. J. Hum. Comput. Interact.* 33, 943–962. doi: 10.1080/10447318.2017.1300750
- Ahmad, M. I., Mubin, O., Shahid, S., and Orlando, J. (2019). Robot's adaptive emotional feedback sustains children's social engagement and promotes their vocabulary learning: a long-term child-robot interaction study. *Adapt. Behav.* 27, 243–266. doi: 10.1177/1059712319844182
- Al Moubayed, S., Beskow, J., Skantze, G., and Granström, B. (2012). "Furhat: a back-projected human-like robot head for multiparty human-machine interaction," in *Cognitive Behavioural Systems*, eds A. Esposito, A. M. Esposito, A. Vinciarelli, R. Hoffmann, and V. C. Müller (Berlin; Heidelberg: Springer), 114–130. doi: 10.1007/978-3-642-34584-5_9
- Al Moubayed, S., and Skantze, G. (2012). "Perception of gaze direction for situated interaction," in *Proceedings of the 4th Workshop on Eye Gaze in Intelligent Human Machine Interaction* (Berlin; Heidelberg: Springer), 1–6. doi: 10.1145/2401836.2401839
- Anzalone, S. M., Boucenna, S., Ivaldi, S., and Chetouani, M. (2015). Evaluating the engagement with social robots. *Int. J. Soc. Robot.* 7, 465–478. doi: 10.1007/s12369-015-0298-7
- Bartneck, C., Kulić, D., Croft, E., and Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *Int. J. Soc. Robot.* 1, 71–81. doi: 10.1007/s12369-008-0001-3
- Baxter, P., Kennedy, J., Vollmer, A.-L., de Greeff, J., and Belpaeme, T. (2014). "Tracking gaze over time in HRI as a proxy for engagement and attribution of social agency," in *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction* (Bielefeld), 126–127. doi: 10.1145/2559636.2559829
- Bless, H., and Greifeneder, R. (2017). "General framework of social cognitive processing," in *Social Cognition*, eds R. Greifeneder, H. Bless, and K. Fiedler (London: Psychology Press), 16–36. doi: 10.4324/9781315648156-2
- Bochkovski, A., Wang, C.-Y., and Liao, H. (2020). Yolov4: Optimal speed and accuracy of object detection. *ArXiv, abs/2004.10934 [Preprint]*.
- Bonferroni, C. (1936). Teoria statistica delle classi e calcolo delle probabilita. *Pubbl. Istit. Super. Sci. Econ. Commer. Firenze* 8, 3–62.
- Carpinella, C. M., Wyman, A. B., Perez, M. A., and Stroessner, S. J. (2017). "The robotic social attributes scale (rosas): development and validation," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction* (Vienna: ACM), 254–262. doi: 10.1145/2909824.3020208
- Castellano, G., Leite, I., Pereira, A., Martinho, C., Paiva, A., and McOwan, P. W. (2010). Affect recognition for interactive companions: challenges and design in real world scenarios. *J. Multim. User Interfaces* 3, 89–98. doi: 10.1007/s12193-009-0033-5
- Castellano, G., Pereira, A., Leite, I., Paiva, A., and McOwan, P. W. (2009). "Detecting user engagement with a robot companion using task and social interaction-based features," in *Proceedings of the 2009 International Conference on Multimodal Interfaces* (Cambridge, MA), 119–126. doi: 10.1145/1647314.1647336
- Choi, B. C., and Pak, A. W. (2005). A catalog of biases in questionnaires. *Prevent. Chron. Dis.* 2, 1–13.
- Corrigan, L. J., Basedow, C., Küster, D., Kappas, A., Peters, C., and Castellano, G. (2015). "Perception matters! engagement in task orientated social robotics," in *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Kobe), 375–380. doi: 10.1109/ROMAN.2015.7333665

- Corrigan, L. J., Peters, C., Castellano, G., Papadopoulos, F., Jones, A., Bhargava, S., et al. (2013). "Social-task engagement: striking a balance between the robot and the task," in *Embodied Commun. Goals Intentions Workshop ICSR*, Vol. 13 (Geneva), 1–7.
- Cui, M., Zhu, M., Lu, X., and Zhu, L. (2019). Implicit perceptions of closeness from the direct eye gaze. *Front. Psychol.* 9:2673. doi: 10.3389/fpsyg.2018.02673
- Del Duchetto, F., Baxter, P., and Hanheide, M. (2020). Are you still with me? Continuous engagement assessment from a robot's point of view. *arXiv preprint arXiv:2001.03515*. doi: 10.3389/frobt.2020.00116
- Exline, R., Gray, D., and Schuette, D. (1965). Visual behavior in a dyad as affected by interview content and sex of respondent. *J. Pers. Soc. Psychol.* 1:201. doi: 10.1037/h0021865
- Goffman, E. (1964). *Behavior in Public Places*. New York, NY: Simon and Schuster.
- Hanson, D. (2006). "Exploring the aesthetic range for humanoid robots," in *Proceedings of the ICCS/CogSci-2006 Long Symposium: Toward Social Mechanisms of Android Science* (Citeseer) (Vancouver, BC), 39–42.
- Haynes, W. (2013). *Holm's Method*. New York, NY: Springer New York, 902. doi: 10.1007/978-1-4419-9863-7_1214
- Holle, H., and Rein, R. (2015). Easydiag: a tool for easy determination of interrater agreement. *Behav. Res. methods* 47, 837–847. doi: 10.3758/s13428-014-0506-7
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* 6, 65–70.
- Kätsyri, J., Förger, K., Mäkäräinen, M., and Takala, T. (2015). A review of empirical evidence on different uncanny valley hypotheses: support for perceptual mismatch as one road to the valley of eeriness. *Front. Psychol.* 6:390. doi: 10.3389/fpsyg.2015.00390
- Kendon, A. (1967). Some functions of gaze-direction in social interaction. *Acta Psychol.* 26, 22–63. doi: 10.1016/0001-6918(67)90005-4
- Kennedy, J., Baxter, P., and Belpaeme, T. (2015). Comparing robot embodiments in a guided discovery learning interaction with children. *Int. J. Soc. Robot.* 7, 293–308. doi: 10.1007/s12369-014-0277-4
- Kleck, R. (1968). Physical stigma and nonverbal cues emitted in face-to-face interaction. *Hum. Relat.* 21, 19–28. doi: 10.1177/001872676802100102
- Kleinke, C. L. (1986). Gaze and eye contact: a research review. *Psychol. Bull.* 100:78. doi: 10.1037/0033-2909.100.1.78
- Kompatsiari, K., Tikhonoff, V., Ciardo, F., Metta, G., and Wykowska, A. (2017). "The importance of mutual gaze in human-robot interaction," in *International Conference on Social Robotics* (Tsukuba: Springer), 443–452. doi: 10.1007/978-3-319-70022-9_44
- Lala, D., Inoue, K., Milhorat, P., and Kawahara, T. (2017). Detection of social signals for recognizing engagement in human-robot interaction. *arXiv [Preprint]*. arXiv:1709.10257.
- Langer, E. J., Fiske, S., Taylor, S. E., and Chanowitz, B. (1976). Stigma, staring, and discomfort: a novel-stimulus hypothesis. *J. Exp. Soc. Psychol.* 12, 451–463. doi: 10.1016/0022-1031(76)90077-9
- Lemaignan, S., Garcia, F., Jacq, A., and Dillenbourg, P. (2016). "From real-time attention assessment to 'with-me-ness' in human-robot interaction," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Christchurch), 157–164. doi: 10.1109/HRI.2016.7451747
- Li, X. A., Florendo, M., Miller, E. L., Ishiguro, H., and Saygin, P. A. (2015). "Robot form and motion influences social attention," in *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Portland, OR), 43–50. doi: 10.1145/2696454.2696478
- Lin, T.-Y., Maire, M., Belongie, S. J., Hays, J., Perona, P., Ramanan, D., et al. (2014). Microsoft coco: Common objects in context. *ArXiv: abs/1405.0312*. doi: 10.1007/978-3-319-10602-1_48
- MacDorman, K. F., Green, R. D., Ho, C.-C., and Koch, C. T. (2009). Too real for comfort? Uncanny responses to computer generated faces. *Comput. Hum. Behav.* 25, 695–710. doi: 10.1016/j.chb.2008.12.026
- Madera, J. M., and Hebl, M. (2019). To look or not to look: acknowledging facial stigmas in the interview to reduce discrimination. *Person. Assess. Decis.* 5:3. doi: 10.25035/pad.2019.02.003
- Madera, J. M., and Hebl, M. R. (2012). Discrimination against facially stigmatized applicants in interviews: an eye-tracking and face-to-face investigation. *J. Appl. Psychol.* 97:317. doi: 10.1037/a0025799
- McDonnell, R., Breidt, M., and Bülthoff, H. H. (2012). Render me Real? Investigating the effect of render style on the perception of animated virtual humans. *ACM Trans. Graph.* 31:91. doi: 10.1145/2185520.2185587
- Minato, T., Shimada, M., Ishiguro, H., and Itakura, S. (2004). "Development of an android robot for studying human-robot interaction," in *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems (Ottawa, ON: Springer)*, 424–434. doi: 10.1007/978-3-540-24677-0_44
- Mori, M., MacDorman, K. F., and Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robot. Autom. Mag.* 19, 98–100. doi: 10.1109/MRA.2012.2192811
- Mumm, J., and Mutlu, B. (2011). "Human-robot proxemics: physical and psychological distancing in human-robot interaction," in *Proceedings of the 6th International Conference on Human-Robot Interaction* (Lausanne), 331–338. doi: 10.1145/1957656.1957786
- O'Brien, H. L., and Toms, E. G. (2010). The development and evaluation of a survey to measure user engagement. *J. Am. Soc. Inform. Sci. Technol.* 61, 50–69. doi: 10.1002/asi.21229
- Paetzel, M., and Castellano, G. (2019). "Let me get to know you better: can interactions help to overcome uncanny feelings?," in *Proceedings of the 7th International Conference on Human-Agent Interaction* (Kyoto), 59–67. doi: 10.1145/3349537.3351894
- Paetzel, M., Castellano, G., Varni, G., Hupont, I., Chetouani, M., and Peters, C. (2018). "The attribution of emotional state - how embodiment features and social traits affect the perception of an artificial agent," in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (Nanjing), 495–502. doi: 10.1109/ROMAN.2018.8525700
- Paetzel, M., and Manuvinakurike, R. (2019). "Can you say more about the location? The development of a pedagogical reference resolution agent," in *Dialog for Good - Workshop on Speech and Language Technology Serving Society (DiGo)* (Stockholm).
- Paetzel, M., Perugia, G., and Castellano, G. (2020). "The persistence of first impressions: The effect of repeated interactions on the perception of a social robot," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (Cambridge, UK), 73–82. doi: 10.1145/3319502.3374786
- Paetzel-Prüsmann, M., Perugia, G., and Castellano, G. (2021). The influence of robot personality on the development of uncanny feelings towards a social robot. *Comput. Hum. Behav.* 120, 1–17. doi: 10.1016/j.chb.2021.106756
- Papadopoulos, F., Küster, D., Corrigan, L. J., Kappas, A., and Castellano, G. (2016). Do relative positions and proxemics affect the engagement in a human-robot collaborative scenario? *Interact. Stud.* 17, 321–347. doi: 10.1075/is.17.3.01pap
- Perugia, G., Boladeras, M. D., Mallofré, A. C., Rauterberg, M., and Barakova, E. (2017). "Modelling engagement in dementia through behaviour. Contribution for socially interactive robotics," in *2017 International Conference on Rehabilitation Robotics (ICORR)* (London), 1112–1117. doi: 10.1109/ICORR.2017.8009398
- Perugia, G., Díaz-Boladeras, M., Catalá-Mallofré, A., Barakova, E. I., and Rauterberg, M. (2020). ENGAGE-DEM: a model of engagement of people with dementia. *IEEE Trans. Affect. Comput.* 1–18. doi: 10.1109/TAFFC.2020.2980275
- Perugia, G., Paetzel-Prüsmann, M., Alanenpää, M., and Castellano, G. (2021). I can see it in your eyes: Gaze towards a robot as an implicit cue of uncanniness and task performance in long-term interactions. *arXiv [Preprint]*. arXiv:2101.05028.
- Perugia, G., Van Berkel, R., Díaz-Boladeras, M., Catalá-Mallofré, A., Rauterberg, M., and Barakova, E. (2018). Understanding engagement in dementia through behavior. The ethnographic and laban-inspired coding system of engagement (ELICSE) and the evidence-based model of engagement-related behavior (EMODEB). *Front. Psychol.* 9:690. doi: 10.3389/fpsyg.2018.00690
- Ray, G. B., and Floyd, K. (2006). Nonverbal expressions of liking and disliking in initial interaction: encoding and decoding perspectives. *South. Commun. J.* 71, 45–65. doi: 10.1080/10417940500503506
- Rosenthal-von der Pütten, A. M., and Krämer, N. C. (2014). How design characteristics of robots determine evaluation and uncanny valley related responses. *Comput. Hum. Behav.* 36, 422–439. doi: 10.1016/j.chb.2014.03.066
- Serholt, S., and Barendregt, W. (2016). "Robots tutoring children: longitudinal evaluation of social engagement in child-robot interaction," in *Proceedings of the 9th Nordic Conference on Human-Computer Interaction* (Gothenburg), 1–10. doi: 10.1145/2971485.2971536

- Sidner, C. L., Lee, C., Kidd, C., Lesh, N., and Rich, C. (2005). Explorations in engagement for humans and robots. *arXiv preprint cs/0507056*. doi: 10.1016/j.artint.2005.03.005
- Smith, M. A., and Wiese, E. (2016). “Look at me now: investigating delayed disengagement for ambiguous human-robot stimuli,” in *International Conference on Social Robotics* (Kansas City, MO: Springer), 950–960. doi: 10.1007/978-3-319-47437-3_93
- Strait, M. K., Vujovic, L., Floerke, V., Scheutz, M., and Urry, H. (2015). “Too much humanness for human-robot interaction: exposure to highly humanlike robots elicits aversive responding in observers,” in *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (Seoul), 3593–3602. doi: 10.1145/2702123.2702415
- Thepsonthorn, C., Ogawa, K.-i., and Miyake, Y. (2021). The exploration of the uncanny valley from the viewpoint of the robot’s nonverbal behaviour. *Int. J. Soc. Robot.* 1–13. doi: 10.1007/s12369-020-0726-w
- Wirth, J. H., Sacco, D. F., Hugenberg, K., and Williams, K. D. (2010). Eye gaze as relational evaluation: averted eye gaze leads to feelings of ostracism and relational devaluation. *Pers. Soc. Psychol. Bull.* 36, 869–882. doi: 10.1177/0146167210370032

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Perugia, Paetzel-Prüßmann, Alanenpää and Castellano. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.