



Biswas, Mriganka and Murray, John (2026) Embodiment and Trust in Human-AI Interaction: a Mixed-Methods Study of Alexa vs. Pepper. International Journal of Social Robotics, 18 (79). ISSN 1875-4791

Version: AAM

Available under license: Creative Commons: Attribution 4.0

Downloaded from: <https://sure.sunderland.ac.uk/id/eprint/20705/>

Which version am I reading (articles, book chapters and books)

Pre-print (or submitted version): the version of an article or chapter before any peer-review. Can be submitted to a publisher or added to a pre-print server.

AAM: Author Accepted Manuscript - The version of the article formally accepted for publication after the peer-review process. It includes the changes made by the author(s) in line with peer-reviewers' comments. It has not undergone the publisher's formatting, type-setting and copyediting.

VOR: Version of Record - The version published on the publisher's website which includes formatting, typesetting and copyediting.

Usage guidelines

Please refer to the usage guidelines at <http://sure.sunderland.ac.uk/policies.html> or alternatively contact sure@sunderland.ac.uk.

Embodiment and Trust in Human-AI Interaction: A Mixed-Methods Study of Alexa vs. Pepper

Dr Biswas M, University of Sunderland, UK, mriganka.biswas@sunderland.ac.uk (corresponding author)

Prof Murray J, University of Huddersfield, UK, john@jcmurray.com

Abstract: As artificial intelligence systems are increasingly deployed in socially meaningful contexts, understanding how their physical form, or embodiment, affects user trust is critical. This mixed-methods study investigates how both physical and perceived embodiment shape user responses to AI agents. Seventy-one participants interacted with two distinct AI systems—Pepper, an embodied humanoid robot, and Alexa, a disembodied voice assistant—while playing a competitive rock-paper-scissors game.

Quantitative analysis revealed that participants rated Pepper significantly higher in both trust and engagement than Alexa. Critically, a mediation analysis demonstrated that this "embodiment advantage" is explained by perception; the user's subjective experience of the AI as being lifelike and physically present fully mediated the relationship between the type of AI and the level of trust.

Qualitative analysis of open-ended responses deepened these findings. Participants described Pepper as more transparent and fair, associating its physical gestures with trustworthiness. Conversely, while Alexa was noted for its predictability, it also evoked privacy concerns and was sometimes perceived as a "black box". The study also found that participants with lower educational backgrounds reported significantly higher perceived embodiment for Pepper, suggesting embodiment effects are moderated by user characteristics. These results underscore that trust is a function of the user's interpretation of embodiment, offering key insights for designing more socially responsive and ethical AI.

Keywords: *Human–AI Interaction, Embodiment, Trust in AI, Perceived Embodiment, CASA Paradigm*

1. Introduction

Artificial intelligence (AI) systems are becoming deeply embedded in daily life, mediating interactions in contexts ranging from customer service to education and healthcare. As these systems become more autonomous and socially integrated, a central challenge emerges: how do humans form trust in AI, and how might design features like embodiment shape this trust? Trust is not a monolithic concept; rather, it encompasses dimensions of reliability, fairness, transparency, and ethical alignment (Mayer, Davis, & Schoorman, 1995). While prior research has extensively examined the functional attributes of AI—such as accuracy, predictability, and efficiency—as drivers of user trust (Hoff & Bashir, 2015), the social and perceptual dimensions of AI systems remain comparatively underexplored. In particular, the role of physical embodiment—the presence of a tangible, expressive form—presents a critical yet understudied design factor in fostering trust and user engagement.

Emerging literature suggests that physical embodiment can enhance social presence, emotional connection, and perceived intelligence in AI systems (Dautenhahn, 2007; Reeves & Nass, 1996). Embodied agents, such as humanoid robots, leverage non-verbal cues like eye gaze, gestures, and

physical proximity to create a more lifelike and socially engaging interaction. By contrast, disembodied AI agents—such as voice assistants—interact solely through auditory channels, often optimised for speed and task efficiency. Yet, while embodied systems have shown promise in collaborative settings (e.g., therapy or tutoring), it remains unclear whether the same advantages hold in competitive contexts, where perceptions of fairness, predictability, and autonomy are paramount. This is a significant gap in the literature, as many AI systems are now being deployed in gaming, negotiation, and decision-making scenarios that involve competition, ambiguity, or strategic behaviour.

Moreover, although previous studies have linked embodiment to higher levels of trust (Bartneck et al., 2009; Glikson & Woolley, 2020), the mechanisms underlying this relationship remain insufficiently understood. Is embodiment itself responsible for increased trust, or is it the *perception* of embodiment—how “real” or agentic the AI seems to the user—that mediates this effect? In other words, does the robot’s physical presence intrinsically enhance trust, or does trust emerge because users perceive the AI as socially aware, autonomous, or emotionally attuned? Addressing this question requires disentangling physical embodiment from perceived embodiment, a distinction that most existing studies overlook.

This paper addresses these gaps by investigating how users trust and engage with two contrasting AI systems: Pepper, a physically embodied humanoid robot, and Alexa, a disembodied voice assistant. Both systems participated in a competitive rock-paper-scissors game with human users, creating a controlled yet dynamic environment where trust, fairness, and strategy could be evaluated. Using a mixed-methods approach, we assessed participants’ trust and engagement levels through both quantitative measures (e.g., Likert-scale questionnaires, mediation modelling) and qualitative responses (e.g., open-ended questions thematically coded).

Importantly, this study goes beyond surface-level comparisons. First, we explore whether perceived embodiment mediates the relationship between physical presence and trust—a novel contribution to embodiment literature that clarifies why users might respond differently to embodied versus disembodied AI. Second, we examine how individual differences (e.g., education level, tech familiarity) moderate these perceptions, uncovering user-specific factors that shape trust in AI. Finally, we position our work within and between key theoretical models: the Capability-Benevolence-Integrity (CBI) framework (Mayer et al., 1995), the Technology Acceptance Model (Davis, 1989), and the CASA paradigm (Reeves & Nass, 1996), using them not just descriptively but analytically to interpret user-AI interactions in a competitive, real-time setting.

Thus, this study makes three distinct contributions. Firstly, it offers a direct experimental comparison of embodied and disembodied AI in a competitive context—an understudied but increasingly relevant domain as AI systems are deployed in environments requiring strategic decision-making and fairness. Secondly, it provides evidence for mediation, demonstrating that users’ trust is shaped not simply by the AI’s physical presence, but by their *perception* of its embodiment—clarifying the psychological mechanism by which trust is formed. Thirdly, it highlights the moderating role of individual traits, particularly education level and technological familiarity, in shaping how users interpret, engage with, and ultimately trust different forms of AI.

Together, these contributions address key blind spots in human-AI interaction research and inform the design of AI systems that are not only functionally competent but also socially intelligent and ethically attuned to users’ expectations.

To guide this investigation, the following research questions are posed:

- **RQ1:** Do users report different levels of trust and engagement when interacting with an embodied AI (Pepper) versus a disembodied AI (Alexa) in a competitive context?
- **RQ2:** Does perceived embodiment mediate the relationship between physical presence and trust?
- **RQ3:** How do individual differences, such as education level or technological familiarity, influence users' perceptions of embodiment and trust?

By exploring these questions through rigorous quantitative and qualitative analyses, this study advances a nuanced understanding of embodiment in human-AI interaction, offering practical and theoretical insights for the future of socially responsive AI design.

2. Literature Review

As artificial intelligence (AI) technologies grow more pervasive, human-AI interaction research has increasingly turned to questions of trust, agency, and social presence. While AI's computational power and accuracy are often praised, trust remains a fundamental determinant of its acceptance and sustained use (Hoff & Bashir, 2015; Siau & Wang, 2018). A particularly salient factor influencing trust in AI is embodiment—the extent to which an AI system is physically present and capable of enacting gestures, facial expressions, or spatial interactions akin to human behaviour (Dautenhahn, 2007). This literature review critically examines the role of embodiment in shaping trust and engagement, highlights the limitations of current research, and establishes the theoretical rationale for the present study's focus on perceived embodiment as a mediating factor.

Trust, broadly conceptualised as a willingness to be vulnerable to another party based on positive expectations of their intentions or behaviour (Mayer et al., 1995), is crucial in the context of AI. Unlike traditional tools, AI systems are often autonomous and opaque in their decision-making processes, necessitating higher levels of user trust for adoption (Lee & See, 2004). Trust in automation research has highlighted three dimensions of trust—competence (capability), integrity, and benevolence—which are central to the Capability-Benevolence-Integrity (CBI) model (Mayer et al., 1995). This model, adapted to AI contexts, has been instrumental in predicting when and why users choose to delegate tasks to AI systems (Hancock et al., 2011). However, functional reliability alone does not guarantee trust. Hoff and Bashir (2015) argue for a distinction between calibrated trust—where trust accurately reflects system performance—and misuse, where users over-trust systems due to perceived social cues. This distinction underscores the relevance of design features that go beyond performance metrics, such as embodiment, which may signal benevolence or integrity through non-verbal cues rather than actual algorithmic behaviour.

Embodiment in AI spans a spectrum from voice-only assistants (e.g., Amazon Alexa) to physically situated humanoid robots (e.g., Pepper, NAO). Embodied AI systems have been associated with increased anthropomorphism (Waytz, Heafner, & Epley, 2014), social presence (Biocca et al., 2003), and engagement (Wiese et al., 2017). Robots capable of gestures, eye contact, and facial expressions tend to elicit stronger relational responses than their disembodied counterparts (Breazeal, 2003; Fong et al., 2003). The Computers Are Social Actors (CASA) paradigm (Reeves & Nass, 1996) and Media Equation Theory posit that people apply social rules to computers and media, treating them as social entities. Accordingly, physical embodiment may invoke social scripts typically reserved for human interaction, such as politeness or empathy. For example, Kidd and Breazeal (2008) showed that people

were more willing to follow advice from an embodied robot than a screen-based one, even when both conveyed identical information.

Yet, the relationship between embodiment and trust is not linear. Studies have shown that excessive anthropomorphism or human-likeness may lead to discomfort, a phenomenon known as the Uncanny Valley (Mori, 1970; MacDorman & Ishiguro, 2006). Robots that appear almost—but not quite—human can elicit feelings of eeriness and distrust. This points to the importance of not merely having a physical form, but designing embodiment that aligns with user expectations and the robot's functional role. Moreover, perceived embodiment—how users subjectively interpret an AI's physical and social presence—may be a more relevant determinant of trust than physical embodiment alone. Heerink et al. (2010) found that user perceptions of a robot's sociability and agency significantly predicted acceptance, even when actual behaviour was held constant. Similarly, de Greeff and Belpaeme (2015) observed that subtle non-verbal behaviours, such as head tilts or gaze cues, can significantly enhance the perceived lifelikeness of a robot, thereby affecting trust.

Beyond the CASA paradigm, CBI model, and TAM, alternative frameworks such as social identity theory (Tajfel & Turner, 1979) may offer additional insights into trust in human-AI interaction. This theory posits that individuals' perceptions of group membership and social categorisation influence their attitudes and behaviours. In the context of AI, users may categorise embodied agents like Pepper as more akin to an in-group (human-like) compared to disembodied agents like Alexa, potentially enhancing trust through perceived similarity. While not the focus of this study, exploring such social-cognitive mechanisms could enrich our understanding of how embodiment shapes relational dynamics, particularly in diverse cultural contexts.

Most studies on embodiment have focused on collaborative contexts, such as education, elder care, and therapy. In these domains, the social capabilities of robots are leveraged to enhance motivation, companionship, or learning (Belpaeme et al., 2018; Wada & Shibata, 2007). Embodiment in such settings has consistently been shown to boost engagement and emotional connection. However, competitive interactions—which introduce elements of strategy, fairness, and potential conflict—have received comparatively little attention. Competitive settings are unique because they foreground users' concerns about manipulation, cheating, or inequality. Trust in such contexts may hinge not only on competence but also on perceived fairness and transparency (Lee & See, 2004; Kraus et al., 2020). Interestingly, while predictability is often viewed positively in task-oriented AI (Hoff & Bashir, 2015), in games it may be interpreted as evidence of manipulation or strategic deception. Some users perceive disembodied agents like Alexa as "too calculated" or "too fast," leading to suspicions of unfair play. Conversely, slower and more visibly deliberative behaviours from embodied robots may be interpreted as more transparent or ethical, even if they reduce efficiency (Eyssel et al., 2012). This dichotomy highlights the trade-off between efficiency and social presence, a theme that remains underexplored in trust research. While Alexa's rapid responses may satisfy users in information retrieval tasks, they might undermine trust in settings requiring relational transparency. Pepper's slower but embodied gestures may, paradoxically, foster more trust in competitive gameplay by making its decision-making appear more comprehensible and relatable.

A key theoretical advance proposed by this study is the mediating role of perceived embodiment in the relationship between physical presence and trust. While physical embodiment offers a tangible foundation for interaction, it is the user's interpretation of this presence—how alive, autonomous, or socially aware the AI appears—that ultimately shapes trust. To our knowledge, few studies have formally tested this mediation. One exception is work by Li (2015), which found that robot anthropomorphism influenced trust only when users reported a high degree of perceived animacy. This

suggests a two-step process: physical features enable perceived embodiment, which in turn affects trust. Applying a mediation model to test this pathway offers a methodological and conceptual contribution that bridges embodiment design with psychological experience.

Another limitation of much existing literature is the treatment of users as a homogeneous group. In reality, trust is shaped by individual differences, including technological literacy, prior AI experience, cultural norms, and educational background (Nomura et al., 2008; Li et al., 2010). For instance, users with higher education levels or greater tech-savviness may be more critical of AI behaviours, exhibiting lower initial trust but greater ability to calibrate that trust over time (Pak et al., 2012). Conversely, users with limited exposure to AI may anthropomorphise embodied agents more readily, leading to inflated trust based on superficial cues. This miscalibration can result in over-trust, especially when embodiment signals benevolence that the system cannot actually fulfil. Recent studies have started to investigate these dynamics. Glikson and Woolley (2020) found that trust in AI varied significantly across demographic and psychological profiles, suggesting the need for adaptive design strategies. However, research that specifically links education level to perceptions of embodiment and trust remains scarce. The present study addresses this gap by examining how users with different educational backgrounds interpret and respond to the embodiment features of Pepper and Alexa.

In addition to design features and user traits, trust in AI is also shaped by broader concerns about privacy and data security. Disembodied agents like Alexa are often cloud-connected and continuously listening, which raises concerns about surveillance and misuse of personal data (Zhou et al., 2021). These concerns can reduce trust, regardless of the agent's performance or behaviour. By contrast, embodied agents are often perceived as more localised and less data-extractive, even when the actual privacy risks may be similar. This perceptual difference suggests that embodiment may inadvertently mitigate privacy concerns by evoking familiarity or domesticity. However, such effects are not universally beneficial. If embodiment leads users to underestimate privacy risks, it may create a false sense of security. Trust, in such cases, becomes ethically problematic, as it is based on misperceptions rather than informed judgments. Thus, the ethical implications of embodiment extend beyond aesthetics or functionality. Designers must consider how physical presence interacts with user assumptions about data use, surveillance, and control. Trust should not only be cultivated but also calibrated—an ideal that calls for transparency and user education alongside engaging design.

Taken together, the existing literature establishes that embodiment plays a critical role in shaping trust and engagement with AI. However, important gaps remain. Firstly, most research has focused on collaborative or assistive contexts, leaving competitive interactions underexplored. Secondly, while physical embodiment has been widely studied, the mediating role of perceived embodiment in fostering trust is under-theorised and rarely tested empirically. Thirdly, little attention has been given to individual differences such as education level, which may influence how users interpret embodiment and calibrate trust. The current study addresses these gaps through a direct comparison of user interactions with Pepper (an embodied humanoid robot) and Alexa (a disembodied voice assistant) in a competitive gameplay context. By combining quantitative methods (including mediation analysis) with qualitative insights from open-ended responses, it explores not just whether embodiment affects trust, but why and for whom. In doing so, it contributes to a more nuanced and ethically informed understanding of trust in human-AI interaction, with implications for both theory and design.

3. Methodology

As mentioned, this study employed a mixed-methods approach, combining a controlled experiment with quantitative and qualitative data collection, to investigate the impact of embodiment on trust and user experience in human-AI interactions. This approach allows for a comprehensive exploration of the

research questions, capturing both objective measures and subjective experiences. A total of 71 participants were recruited for this study through advertisements and personal invitations. The sample consisted of adults with varying levels of technological experience and educational backgrounds. The demographic characteristics of the participants were as follows:

- **Age:** The age of participants ranged from 18 to 65 years, with a mean age of 32.5 years (SD = 10.2).
- **Gender:** 42 participants (59.2%) identified as female, and 29 participants (40.8%) identified as male.
- **Ethnicity:** The sample predominantly consisted of White British participants (83.1%), followed by Asian (8.5%), Black (4.2%), and other ethnicities (4.2%).
- **Education Level:** 35 participants (49.3%) had completed high school, 22 participants (31%) held a bachelor's degree, 10 participants (14.1%) had a master's degree, and 4 participants (5.6%) had a doctoral degree.

The experimental sessions were conducted individually in a controlled laboratory environment at the University of Sunderland, UK to minimize distractions and ensure focused interactions with the AI entities. The core of the experiment involved participants engaging in multiple rounds of a rock-paper-scissors game against two distinct AI entities:

- **Alexa:** A voice-based AI assistant, representing a disembodied AI, as depicted in Figure 1 (Figure 1, above).
- **Pepper:** A humanoid robot with physical embodiment, as depicted in Figure 1 (Figure 1, below).

The rock-paper-scissors game was chosen for its simplicity, universal familiarity, and inherent strategic decision-making element, ensuring that all participants could easily understand and engage in the gameplay. This game provided a controlled environment for a direct comparison of interactions with an embodied and a disembodied AI, as both entities were tasked with the same objectives within the same game structure.

Participants engaged in separate interactions with Pepper and Alexa, with the order of interaction counterbalanced to mitigate potential order effects. The specifics of the interaction with each AI are detailed in Table 1.

Table 1: Rock-Paper-Scissor design for Alexa and Pepper

	Alexa	Pepper
Interaction	Participants played against Alexa using voice commands. Alexa was set up on a smart speaker in the same room with Pepper, ensuring its responses were clearly audible.	Participants played rock-paper-scissors against Pepper in a face-to-face setting. Pepper was positioned at a table across from the participant.
Game Play	Gameplay was autonomous.	Gameplay was autonomous.
Prompts, Actions & Responses	Interaction was purely verbal. It would initiate each round by prompting the participant with, "What do you choose? Rock, paper, or scissors?" The participant would then state their choice aloud. Alexa subsequently announces its own choice and declare the winner of the round based on the standard rules of the game.	Pepper's behaviour was programmed based on its abilities, e.g., verbal interactions, hand gestures. Pepper shakes its hand three times then asked participants about their choices. Participants then reveals their choice, and after that Pepper states its choices and announces the outcome of the round (win, lose, draw) based on its own choice and the participant's action.

Experimenter's Role	The experimenter's role was limited to ensuring the participant's comfort and the technical setup's functionality.	The experimenter's role was limited to ensuring the participant's comfort and the technical setup's functionality.
----------------------------	--	--

Following the game sessions, participants completed a questionnaire designed to assess their perceptions of trust and embodiment toward both Pepper and Alexa. The questionnaire included the following:

- Demographic questions: Participants provided information about their age, gender, ethnicity, education level, and residency.
- Likert-scale items: These items employed 5-point scales (ranging from "strongly disagree" to "strongly agree") to measure various aspects of:
 - Trust: The trust scale consisted of 10 items adapted from the Trust in Automation scale (Hoff & Bashir, 2015). The scale demonstrated high internal consistency in this study




Figure 1: Participant interacting with Alexa (above) and with Pepper (below)

(Cronbach's $\alpha = .89$). Sample items include: *"To what extent do you feel comfortable relying on this AI to make decisions for you?"*; *"How much do you believe this AI operates in a fair and unbiased manner?"*

- Embodiment: The embodiment scale was composed of 8 items developed for this study. The scale showed good internal consistency (Cronbach's $\alpha = .85$). Sample items include: *"To what extent does this AI seem to have a physical presence in the world?"*; *"How much does this AI feel lifelike and capable of independent action?"*
- Open-ended questions: These questions were designed to elicit more detailed and nuanced insights into participants' experiences interacting with the AI entities. Examples of open-ended questions include: *"Describe your overall experience playing against Alexa compared to Pepper."*; *"Did you feel that either AI had an advantage in the game? Why or why not?"*; *"How did the different ways of interacting with Alexa and Pepper affect your experience?"*

To explore the underlying structure of the trust and embodiment scales and to confirm their unidimensionality, exploratory factor analyses were conducted. The results supported a single-factor solution for each scale, confirming that the items were measuring a single underlying construct.

Qualitative data from the open-ended responses underwent thematic analysis using an inductive approach, facilitated by NVivo software. This involved a systematic and iterative process of identifying, analysing, and reporting patterns within the data. To enhance the rigor of the qualitative analysis, an intercoder reliability check was performed. A second coder independently coded 20% of the responses, and the level of agreement between the two coders was calculated using Cohen's kappa ($\kappa = .82$), indicating substantial agreement. The assignment of multiple codes to some responses reflects the interconnectedness of the themes and the richness of the participants' experiences. Examples of the coding process and theme development are provided in Table 2.

Table 2: Examples of Coding, Theme Development and Quotes

Theme	Description	Frequency (% of responses)	Quote
Lifelikeness and Social Presence	Participants perceived Pepper as more human-like, engaging, and socially present due to its physical embodiment and gestures.	68	“Pepper felt like playing with a real person—it was more engaging than a machine.”
Transparency and Fairness	Pepper’s visible deliberation (e.g., gestures, pauses) was associated with fairness and trustworthiness, unlike Alexa’s rapid responses.	54	“Pepper seemed to think before choosing, which made it feel fairer than Alexa.”
Privacy Concerns	Alexa was associated with surveillance and data misuse, while Pepper was perceived as less intrusive.	42	“Alexa feels like it’s always listening; I don’t trust where my voice goes.”
Familiarity and Task-Fit	Some preferred Alexa for functional tasks due to daily use, but Pepper was favoured for relational gameplay.	35	“I’d choose Alexa for quick tasks, but Pepper made the game more fun.”
Unexpected: Over-Anthropomorphism	Some participants attributed excessive human-like intentions to Pepper, leading to unrealistic expectations (e.g., assuming it could empathise).	15	“I felt Pepper understood my strategy, like it was reading my mind, which was a bit odd.”

Prior to their participation, all participants provided informed consent. The study design was approved by the University of Sunderland Ethics Committee, ensuring adherence to ethical guidelines for research involving human subjects. Participant confidentiality was maintained throughout the study, and data were anonymized during analysis and reporting.

4. Result Analysis & Discussions

The data gathered from the 71 participants who interacted with both the embodied AI, Pepper, and the disembodied AI, Alexa, in a competitive rock-paper-scissors game, yielded a rich tapestry of quantitative and qualitative insights. This section will first present the quantitative findings related to trust, engagement, perceived embodiment, and the influence of individual differences. Subsequently, it will delve into the qualitative themes that emerged from participants' open-ended responses, providing

depth and context. Finally, these quantitative and qualitative results will be interwoven in a comprehensive discussion, interpreting their collective meaning, connecting them to the theoretical frameworks outlined previously, and considering their implications for understanding and designing future human-AI interactions.

Quantitative analysis consistently revealed a significant advantage for the embodied AI, Pepper, across key perceptual measures. Participants reported significantly higher levels of trust in Pepper ($M = 4.18$, $SD = 0.67$) compared to Alexa ($M = 3.87$, $SD = 0.79$). This difference, confirmed by paired sample t -tests with a p -value $< .001$, directly addresses the first research question and confirms that physical embodiment—operationalised through Pepper’s humanoid form—does elicit greater user trust than a disembodied AI counterpart. From a theoretical standpoint, this outcome affirms the predictions of the Computers Are Social Actors (CASA) paradigm (Reeves & Nass, 1996) and supports the Capability-Benevolence-Integrity (CBI) model of trust (Mayer et al., 1995), in which visible behaviours and inferred social presence contribute to users’ evaluations of benevolence and integrity.

Similarly, Pepper outperformed Alexa on perceived user engagement (Pepper: $M = 4.30$, $SD = 0.72$; Alexa: $M = 3.65$, $SD = 0.81$), reinforcing the view that embodiment enhances the richness and meaningfulness of interaction (see Figure 2). Engagement, in this context, was not merely operationalised as attentiveness or enjoyment, but as a user’s sense of relational involvement in the task. The physicality and gestures of Pepper—its ability to mimic “thinking,” perform hand movements, and maintain simulated eye contact—likely contributed to these engagement effects by facilitating social presence and perceived intentionality. These results extend findings from previous studies on social robots in educational or therapeutic contexts (e.g., Wada & Shibata, 2007; Belpaeme et al., 2018) into the less explored domain of competitive HAI, where strategic behaviours and fairness concerns play a larger role.

Additionally, perceived embodiment scores reflected the intuitive yet important distinction between a humanoid robot and a voice-based AI. Participants rated Pepper’s embodiment significantly higher than Alexa’s, capturing the extent to which users attributed physical presence, autonomy, and lifelikeness to each system. The clear differentiation in embodiment perception—despite both systems following identical task scripts—demonstrates that it is not functionality alone that shapes user interpretation, but the presentation and modality of the interaction. This directly supports the rationale for exploring perceived embodiment as a mediator, rather than treating physical form as a direct predictor of trust.

To investigate the psychological mechanism by which embodiment influences trust, a mediation analysis was conducted using the Baron and Kenny (1986) approach. The hypothesis, grounded in the CASA paradigm and extended by literature on social presence and anthropomorphism, was that perceived embodiment would mediate the effect of AI type (embodied vs. disembodied) on user trust. In essence, the study sought to test whether it is not merely the robot’s physical form, but the user’s subjective experience of that form as socially and cognitively “real,” that accounts for increased trust.

The mediation model showed a strong total effect of AI type on trust difference ($\beta = 0.393$, $p < .001$), a significant effect of AI type on perceived embodiment difference ($\beta = 0.542$, $p < .001$), and a significant effect of perceived embodiment difference on trust difference ($\beta = 0.508$, $p < .001$). Crucially, when the mediator was included, the direct effect of AI type on trust difference became non-significant ($\beta = 0.117$, $p = .375$), confirming full mediation. This model is illustrated conceptually in Figure 3.




Figure 2: Comparison of Mean Trust and Engagement Scores for Pepper and Alexa

These results address the second research question and offer a substantial contribution to human-AI interaction literature by empirically confirming that embodiment functions psychologically, not just materially. While physical presence may invite social responses, it is the interpretation of that presence as intentional, relatable, or agentic that enables trust. This mediating process supports and extends the theoretical claims made by Heerink et al. (2010) and Waytz et al. (2014) that anthropomorphism and perceived animacy are central to users’ affective evaluations of AI.




Figure 3: Mediation Model: AI Type, Perceived Embodiment, and Trust

In addition to mediation analysis, a factor analysis was conducted to explore whether trust and embodiment scores reflected a shared underlying construct. The scree plot, as illustrated in Figure 4, revealed a dominant factor accounting for 43.26% of total variance, with high loadings from Trust_Pepper (0.823) and Embodiment_Pepper (0.662). This reinforces the idea that user perceptions of trust and embodiment are not fully distinct but coalesce around a latent experiential dimension—perhaps best understood as “relational engagement.” These results echo social presence theory

(Gunawardena, 1995), where perceived co-presence significantly influences relational and affective outcomes in mediated interaction.

Turning to individual differences, one of the most striking findings was the influence of education level on perceived embodiment. Participants with lower formal education reported significantly higher perceived embodiment scores for Pepper than those with higher educational qualifications. This suggests that the cues that signify “presence” or “realness” are differentially interpreted across educational backgrounds—possibly due to differences in exposure to technology, cognitive framing, or default anthropomorphism tendencies. While speculative, one interpretation is that users with less formal education may find tangible, embodied cues more intuitively meaningful, while highly educated users might engage in more sceptical or analytical appraisals of AI behaviour.

This finding partially addresses the third research question and opens important ethical questions. Designers of AI systems may need to consider how embodiment cues could be differently persuasive—or misleading—to different demographic groups. Trust that is built on affective or superficial cues, without an understanding of underlying capabilities or limitations, may create a false sense of security. Thus, embodiment design must be paired with strategies for transparency and user education, especially for vulnerable or less technologically literate populations.

Other demographic variables, such as gender and self-reported tech-savviness, were not significantly associated with differences in trust or perceived embodiment. Interestingly, while tech-savviness did not appear to predict trust, there was a moderate positive correlation between perceived fairness and trust in Alexa ($r = 0.45$, $p = .012$), suggesting that in the absence of physical cues, consistent and understandable behaviour becomes the primary basis for trust. This finding is especially relevant to disembodied AI, where predictability and perceived fairness may serve as compensatory signals for the lack of relational warmth.




Figure 4. Scree Plot of Factor Analysis.

Qualitative data enriched the quantitative findings, revealing distinct user perceptions of Pepper and Alexa. Participants frequently described Pepper as lifelike and socially present, with comments like “It was like playing with a real person” highlighting its ability to evoke human-like engagement. This aligns with the CASA paradigm, suggesting embodied cues trigger social scripts. Conversely, Alexa’s rapid responses were often perceived as “cold” or “mechanical,” with some suspecting unfairness due to its speed, underscoring a trade-off between efficiency and relational trust. Privacy concerns further

differentiated the systems: Alexa was seen as intrusive (“always listening”), while Pepper felt “local” and less threatening, despite identical data configurations. An unexpected finding was over-anthropomorphism, where some attributed excessive intentions to Pepper, such as “reading my mind,” suggesting potential miscalibration of trust. These themes, detailed in Table 2, highlight embodiment’s role in shaping trust, fairness, and emotional connection in competitive interactions.

A dominant theme was the perception of Pepper as more lifelike and socially present. Participants repeatedly described the robot using terms such as “real,” “engaging,” “thoughtful,” and “relatable.” One participant remarked, “Pepper’s physical embodiment made it seem more intelligent,” while another said, “It was like having someone physically to interact with—it felt more like playing with a person than a machine.” These reflections resonate strongly with the CASA paradigm and Social Presence Theory, both of which posit that embodied agents trigger social scripts and interpersonal evaluations akin to human-to-human interaction.

Many participants also associated Pepper’s embodiment with enhanced trustworthiness. The ability to see and anticipate its actions—gestures, head tilts, or pauses—was often interpreted as a form of transparency. Several participants noted that Pepper seemed to “think” before responding, which paradoxically made it appear more ethical or fair. Alexa, on the other hand, was often perceived as too fast, even “cold” or “mechanical,” leading some to suspect it of “cheating” or being too calculating. This distinction is especially important in the context of competitive AI, where fairness and predictability become focal points of judgement. The robot’s visible deliberation fostered a stronger sense of procedural justice, even if it was technically slower or less efficient than its disembodied counterpart.

This tension between predictability and perceived fairness offers valuable insights for AI design. While predictable behaviour is often associated with trust in task-based contexts (Hoff & Bashir, 2015), in socially nuanced or competitive settings, it may backfire if it is seen as overly algorithmic or impersonal. Embodiment, by making internal processes appear more transparent, may serve as a powerful antidote to this form of suspicion. From the perspective of the CBI model, Pepper’s physical presence may have enhanced perceptions of benevolence and integrity—dimensions often overlooked when AI is evaluated solely in terms of technical performance.

Distribution of Participant Privacy Concerns by AI Type




Figure 5. Distribution of Participant Privacy Concerns by AI Type

Privacy emerged as another critical axis along which the two systems were evaluated differently. Participants expressed notably more privacy-related concerns regarding Alexa as shown in Figure 5.

Statements such as “Alexa feels like it’s always listening” or “I don’t trust where my voice goes with Alexa” were common. By contrast, Pepper was more frequently described as “local,” “self-contained,” or “less intrusive,” despite the fact that both systems were disconnected from cloud-based processing during the experiment. This perception gap illustrates how embodiment may serve as a heuristic for data security and surveillance potential—a cognitive shortcut that links physical tangibility to assumed privacy. From a design ethics standpoint, this raises concerns. While embodiment may ease anxiety, it can also mislead users about the actual risks and boundaries of data collection.

A final emergent theme related to familiarity and task fit. Some participants expressed a preference for Alexa in non-social or functional domains, citing their everyday experience using it for information retrieval or home control. This aligns with the Technology Acceptance Model’s emphasis on perceived ease of use. However, for interactive or relational tasks such as gameplay, participants overwhelmingly favoured Pepper, noting that it made the interaction more “fun,” “enjoyable,” or “authentic.” This distinction underscores the importance of contextual trust—a recognition that trust in AI is not monolithic but contingent on the nature of the task and the expectations users bring to it.

Together, the qualitative findings converge with the quantitative evidence to paint a coherent and theoretically grounded picture. The study robustly answers its first research question: embodiment, operationalised through Pepper’s physical presence and movements, significantly enhances trust and engagement. It also addresses the second question by showing that this effect operates largely through perception—users trust what they perceive as socially aware and lifelike, not necessarily what is more competent or efficient. Finally, by identifying the moderating role of education and highlighting the absence of tech-savviness effects, the study provides preliminary but compelling evidence for the role of individual differences in shaping HAI perceptions.

An additional consideration is the potential influence of confounding variables, notably participants’ prior familiarity with Alexa and the duration of interactions. Alexa, as a widely used consumer product, may have elicited preconceived notions of reliability or privacy concerns, potentially skewing trust and engagement scores compared to Pepper, a novel research robot. Similarly, the brief duration of the rock-paper-scissors interactions may have amplified the novelty effect of Pepper’s embodiment while limiting opportunities to assess Alexa’s relational capabilities. These factors suggest that embodiment effects may be partly context-dependent, shaped by users’ prior experiences and interaction length. Future studies should control for familiarity by including baseline measures of AI exposure and extend interaction durations to explore how trust evolves over time, ensuring a more robust understanding of embodiment’s role in human-AI interaction.

Taken together, these findings make a meaningful contribution to ongoing theoretical debates and practical design considerations in human-AI interaction. First, they refine and extend the CASA paradigm, demonstrating that the mere presence of a humanoid body is insufficient—it must be *perceived* as socially aware and agentic to elicit social responses such as trust, empathy, or engagement. The current study advances this paradigm by demonstrating not only that social cues matter, but that their interpretation is what shapes user experience. Thus, embodiment becomes a psychological phenomenon, not just a technological feature.

Second, the findings strongly reinforce and add complexity to the Capability-Benevolence-Integrity (CBI) framework. Traditional models have often assumed that trust is shaped predominantly by observable performance, yet our study shows that relational signals—such as physical expressiveness and perceived presence—play an equally, if not more, important role in cultivating perceived benevolence and integrity. Pepper’s embodiment led participants to judge it as more trustworthy, not because it won more games, but because it felt more *fair*, *transparent*, and *engaged* in the interaction.

Third, the results nuance the Technology Acceptance Model (TAM) by illustrating how *usefulness* and *ease of use* are interpreted differently depending on context. Alexa, despite being a familiar and efficient assistant, was not trusted as much in a competitive or emotionally evaluative task. This finding has implications for the deployment of AI in settings such as education, therapy, or negotiation, where relational richness and user comfort are paramount. In such contexts, embodiment may not only increase engagement but also redefine what users perceive as useful or trustworthy.

Crucially, the study offers a forward-looking reflection on the ethics of AI design. The differential interpretations based on education level serve as a reminder that AI systems are not perceived uniformly across users. The same design feature—such as a humanoid robot—may inspire awe and trust in one user, and scepticism or discomfort in another. This underscores the need for inclusive design practices and greater awareness of cognitive diversity in AI perception. Furthermore, while embodiment can foster trust, it can also obscure real privacy risks, as seen in users' misplaced comfort with Pepper over Alexa despite equivalent backend data configurations. Ethical design must therefore ensure that social cues do not override informed judgement, especially among less technically literate populations.

In sum, this study contributes a multi-dimensional understanding of how embodiment functions in AI-mediated interaction—empirically, psychologically, and socially. It answers foundational questions about trust, engagement, and user interpretation, while also opening new lines of inquiry into individual differences, contextual appropriateness, and ethical AI deployment. By integrating quantitative analysis, qualitative insight, and critical theoretical reflection, it lays a robust foundation for future work in the field. Further research should explore how these patterns evolve over time, whether trust in embodied AI sustains through repeated interaction, and how culture, language, and emotion interact with physical presence to shape relational dynamics. As AI systems continue to proliferate in both public and private domains, understanding the subtle interplay between embodiment, perception, and trust will be essential—not just for system efficiency, but for human dignity, autonomy, and well-being.

5. Conclusion

This study set out to explore the complex interplay between physical embodiment, perceived embodiment, and trust in human-AI interactions, using a controlled yet ecologically valid competitive game scenario to elicit rich behavioural and perceptual data. By comparing an embodied humanoid robot (Pepper) with a disembodied voice assistant (Alexa), and by employing a mixed-methods approach that integrated quantitative modelling and qualitative thematic analysis, the research offers a multidimensional account of how embodiment shapes user experience, trust, and engagement with artificial intelligence.

A central and novel contribution of this study lies in its demonstration that physical embodiment alone is not the decisive factor in fostering trust—rather, it is the *user's perception* of embodiment that mediates this relationship. This distinction clarifies a previously blurred theoretical space in the literature. Through mediation analysis, the study provides empirical support for the claim that trust is not an automatic by-product of humanoid form, but a consequence of how human users interpret and make meaning of that form within social and task contexts. This insight reinforces the idea that embodiment is as much a psychological and communicative construct as it is a physical one.

Moreover, the study situates this insight within a broader context of individual user differences. Findings that education level influenced perceptions of embodiment add a layer of complexity to the design of AI systems: trust is not universally experienced, and design cues may not function identically

across different populations. While many models of AI adoption treat the user base as a homogenous group, this study shows that demographic, cognitive, and perceptual diversity must be taken seriously when developing AI that is not only efficient but also ethically sound and socially inclusive.

The implications extend beyond academic theory. As AI systems are increasingly integrated into public-facing roles—customer service agents, healthcare companions, educational tutors—the question of how they are perceived, trusted, and socially engaged becomes a design issue of ethical consequence. The study's findings suggest that in tasks where relational presence, fairness, and interpretability matter, embodied AI agents have a clear advantage—but only when their design supports believable and appropriately human-like interaction. Conversely, disembodied agents may still be preferred for functional efficiency, provided they signal consistency, transparency, and ethical handling of data.

The findings from this study offer several critical implications for the design of trustworthy and engaging AI systems. First, designers should recognise that perceived embodiment, rather than physical form alone, is the primary driver of user trust. This suggests that trust can be enhanced not only through humanoid design, but through behavioural cues that foster social presence—such as timing, responsiveness, gestures, and perceived deliberation. Second, trust is context-sensitive. In competitive settings, users valued fairness, transparency, and social cues more than efficiency. While Alexa's predictability was appreciated, it often translated into mechanical coldness, undermining relational trust. Designers should therefore calibrate the agent's responsiveness and strategy to the interaction context—making AI more legible, particularly in adversarial or emotionally evaluative tasks. Third, the study highlights the need to account for individual differences. For instance, participants with lower educational backgrounds were more influenced by physical embodiment cues. This raises important questions around cognitive vulnerability and the ethical use of persuasive design. Future AI systems must balance engagement with transparency and fairness, avoiding manipulative cues that may disproportionately affect less technically literate users. Finally, privacy perceptions were strongly influenced by embodiment. Users associated disembodied, cloud-connected AI with greater surveillance risk. Designers of voice-only systems must work to visualise or explain privacy boundaries more clearly. By contrast, embodied agents should not rely solely on their physical charm to offset concerns; instead, embodiment should be coupled with explicit indicators of data safety and consent. Taken together, these implications suggest that the future of ethical, trustworthy AI lies in context-aware, user-sensitive, and socially legible design—a principle especially vital in competitive, emotionally charged, or trust-critical domains.

The study's findings, while robust, are subject to several limitations that warrant consideration. As noted, the sample size of 71 participants, though sufficient for statistical analysis, was drawn from a specific cultural and institutional context, predominantly White British, which may constrain the generalisability of results across diverse populations. Similarly, the controlled laboratory setting, while facilitating precise comparisons, may not fully capture the dynamics of real-world human-AI interactions, where environmental factors and prolonged engagement could alter trust perceptions. Beyond these, the study's reliance on a single competitive task—rock-paper-scissors—introduces potential task specificity. While this game provided a controlled and engaging context, its simplicity and short duration may not reflect the complexities of other competitive or relational scenarios, such as negotiation or collaborative decision-making, where trust dynamics might differ. Additionally, participants' prior familiarity with Alexa, a widely used consumer product, compared to Pepper, a less familiar research robot, may have influenced perceptions of trust and engagement, potentially confounding the effects of embodiment. Cultural influences were also not explored, despite the likelihood that norms around anthropomorphism and technology acceptance vary globally. Future research should address these limitations by employing larger, more diverse samples, longitudinal

designs, and varied interaction contexts to better understand the interplay of embodiment and trust in human-AI interaction.

Looking ahead, several research directions emerge. First, longitudinal studies could track how trust in embodied and disembodied AI evolves over time, especially as novelty effects wane. Second, cross-cultural replications would be invaluable in understanding how embodiment is interpreted differently across social and cultural contexts. Third, as AI systems become multimodal—combining voice, gesture, facial expression, and adaptive behaviour—it will be essential to study how these cues interact to create coherent, trustworthy experiences. Finally, there is a pressing need for participatory design approaches that include users from diverse backgrounds in the co-development of AI agents, ensuring that trust is not simply engineered but co-constructed.

In sum, this study affirms the critical importance of embodiment—not merely as a structural feature of AI, but as a relational, interpretive, and psychological catalyst for trust. It reveals the subtle dynamics by which physical presence can either bridge or widen the gap between humans and machines. As we move toward a future in which AI is increasingly embedded in the social fabric, it is this nuanced understanding of trust—not just as a metric, but as a lived experience—that will determine whether artificial intelligence serves as a functional tool or a genuine partner in human life.

Acknowledgements: The authors would like to thank the 71 participants of this study for their time and valuable contributions. We are also grateful to our colleagues at the University of Sunderland who provided feedback on early drafts of this manuscript.

Statements and Declarations:

The authors declare they have no financial or non-financial competing interests to disclose.

The study design was approved by the University of Sunderland Ethics Committee, ensuring adherence to ethical guidelines for research involving human subjects.

The datasets and analysed during the current study are not publicly available but are available from the corresponding author on reasonable request.

REFERENCES

- Baron, R. M., & Kenny, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6), 1173–1182. <https://doi.org/10.1037/0022-3514.51.6.1173>
- Bartneck, C., Kulic, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1), 71–81. <https://doi.org/10.1007/s12369-008-0002-3>
- Belpaeme, T., Kennedy, J., Baxter, P., Vogt, P., Krahmer, E., Kopp, S., & others. (2018). L2TOR—Second language tutoring using social robots. *Advanced Robotics*, 32(10), 1–17. <https://doi.org/10.1080/01691864.2018.1488964>
- Biocca, F., Harms, C., & Burgoon, J. K. (2003). Toward a more robust theory and measure of social presence: Review and suggested criteria. *Presence: Teleoperators and Virtual Environments*, 12(5), 456–480. <https://doi.org/10.1162/105474603322761270>
- Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies*, 59(1–2), 119–155. [https://doi.org/10.1016/S1071-5819\(03\)00018-1](https://doi.org/10.1016/S1071-5819(03)00018-1)
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>

Dautenhahn, K. (2007). Socially intelligent robots: Dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480), 679–704. <https://doi.org/10.1098/rstb.2006.2004>

de Greeff, J., & Belpaeme, T. (2015). Why robots should be social: Enhancing machine learning through social human-robot interaction. *PLOS ONE*, 10(9), e0138061. <https://doi.org/10.1371/journal.pone.0138061>

Eyssel, F., Kuchenbrandt, D., Bobinger, S., de Ruiter, L., & Hegel, F. (2012). ‘If you sound like me, you must be more human’: On the interplay of vocal similarity and human-likeness in trust towards social robots. *Proceedings of the 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 147–148. <https://doi.org/10.1145/2157689.2157730>

Fong, T., Nourbakhsh, I., & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3–4), 143–166. [https://doi.org/10.1016/S0921-8890\(02\)00372-X](https://doi.org/10.1016/S0921-8890(02)00372-X)

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y. C., de Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. *Human Factors*, 53(5), 517–527. <https://doi.org/10.1177/0018720811417254>

Heerink, M., Kröse, B., Evers, V., & Wielinga, B. (2010). Assessing acceptance of assistive social agent technology by older adults: The Almere model. *International Journal of Social Robotics*, 2(4), 361–375. <https://doi.org/10.1007/s12369-010-0068-5>

Hoff, K. A., & Bashir, M. (2015). Trust in automation: Integrating empirical evidence on factors that influence trust. *Human Factors*, 57(3), 407–434. <https://doi.org/10.1177/0018720814547570>

Kidd, C. D., & Breazeal, C. (2008). Robots at home: Understanding long-term human-robot interaction. *Proceedings of the 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3230–3235. <https://doi.org/10.1109/IROS.2008.4651113>

Kraus, J., Scholz, D., & Pleissner, J. (2020). Effects of longitudinal human-robot interaction on trust development in the industrial context. *Frontiers in Robotics and AI*, 7, 545. <https://doi.org/10.3389/frobt.2020.545584>

Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Li, D. (2015). Effect of anthropomorphism on trust in human-robot interaction. *International Journal of Social Robotics*, 7(5), 657–669. <https://doi.org/10.1007/s12369-015-0309-8>

Li, D., Rau, P. P., & Li, Y. (2010). A cross-cultural study: Effect of robot appearance and task. *International Journal of Social Robotics*, 2(2), 175–186. <https://doi.org/10.1007/s12369-010-0056-9>

MacDorman, K. F., & Ishiguro, H. (2006). The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoperators and Virtual Environments*, 15(4), 393–407. <https://doi.org/10.1162/pres.15.4.393>

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709–734. <https://doi.org/10.5465/amr.1995.9508080335>

Mori, M. (2012). The uncanny valley (K. F. MacDorman & N. Kageki, Trans.). *IEEE Robotics & Automation Magazine*, 19(2), 98–100. <https://doi.org/10.1109/MRA.2012.2192811> (Original work published 1970 in *Energy*, 7(4), 33–35)

Nomura, T., Suzuki, T., Kanda, T., & Kato, K. (2008). Measurement of negative attitudes toward robots. *Interaction Studies*, 9(3), 437–454. <https://doi.org/10.1075/is.9.3.11nom>

- Pak, R., Fink, N., Price, M. M., Bass, B., & Sturre, L. (2012). Decision support aids with anthropomorphic characteristics influence trust and performance in younger and older adults. *Ergonomics*, 55(9), 1059–1072. <https://doi.org/10.1080/00140139.2012.691554>
- Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press.
- Short, J., Williams, E., & Christie, B. (1976). *The social psychology of telecommunications*. Wiley.
- Siau, K., & Wang, W. (2018). Building trust in artificial intelligence, machine learning, and robotics. *Cutter Business Technology Journal*, 31(2), 47–53.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Brooks/Cole.
- Wada, K., & Shibata, T. (2007). Living with seal robots—Its sociopsychological and physiological influences on the elderly at a care house. *IEEE Transactions on Robotics*, 23(5), 972–980. <https://doi.org/10.1109/TRO.2007.906261>
- Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113–117. <https://doi.org/10.1016/j.jesp.2014.01.005>
- Zhou, Y., Zhang, J., & Wang, W. (2021). Privacy concerns in voice assistant usage: A multi-dimensional perspective. *Computers in Human Behavior*, 123, 106873. <https://doi.org/10.1016/j.chb.2021.106873>