

Exploring Interactions Between Trust, Anthropomorphism, and Relationship Development in Voice Assistants

WILLIAM SEYMOUR*, University of Oxford, UK

MAX VAN KLEEK, University of Oxford, UK

Modern conversational agents such as Alexa and Google Assistant represent significant progress in speech recognition, natural language processing, and speech synthesis. But as these agents have grown more realistic, concerns have been raised over how their social nature might unconsciously shape our interactions with them. Through a survey of 500 voice assistant users, we explore whether users' relationships with their voice assistants can be quantified using the same metrics as social, interpersonal relationships; as well as if this correlates with how much they trust their devices and the extent to which they anthropomorphise them. Using Knapp's staircase model of human relationships, we find that not only can human-device interactions be modelled in this way, but also that relationship development with voice assistants correlates with increased trust and anthropomorphism.

CCS Concepts: • **Human-centered computing** → **Empirical studies in interaction design**; **Natural language interfaces**.

Additional Key Words and Phrases: voice assistants, trust, anthropomorphism, relationship development

ACM Reference Format:

William Seymour and Max Van Kleek. 2021. Exploring Interactions Between Trust, Anthropomorphism, and Relationship Development in Voice Assistants. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 371 (October 2021), 16 pages. <https://doi.org/10.1145/3479515>

1 INTRODUCTION

In the 1927 silent film *Metropolis*, the inventor Rotwang creates a 'maschinenmensch' to resurrect his dead lover [42]. One of the first robots to ever appear in film, its human appearance is used to manipulate relationships amongst the main characters and start an uprising. Several decades later, Asimov coined the term robotics and predicted the social and emotional challenges of interacting with realistic 'humaniform' robots [5]. More recent films such as *Her* and *Ex Machina* have continued to explore relationships with machines that look and/or sound human in a modern context [29, 35].

In the 1990's, Nass et al. showed that we are unable to engage rationally with voice interfaces; speech is so tightly integrated with many interpersonal functions in the brain that we cannot help but respond socially to devices with voices [56, 57, 64]. Weizenbaum had already said of ELIZA "how easy it is to create and maintain the illusion of understanding [...]. A certain danger lurks there" [73], and Rosalind Picard was cognisant of the capability of affective computers to mislead [58]. Despite this, it was taken for granted in the early days of conversational agents that human-like, embodied agents were the "logical extension of the metaphor of human-computer interaction as a conversation" [16]. It seemed obvious to make use of communication skills that people already possessed, as well as the fact that natural language has built in redundancy and

*Now at King's College London: william.1.seymour@kcl.ac.uk, Bush House, 30 Aldwych, London, WC2B 4BG

Authors' addresses: William Seymour, william.seymour@cs.ox.ac.uk, University of Oxford, Parks Road, Oxford, UK, OX1 3QD; Max Van Kleek, max.van.kleek@cs.ox.ac.uk, University of Oxford, Parks Road, Oxford, UK, OX1 3QD.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the ACM on Human-Computer Interaction*, <https://doi.org/10.1145/3479515>.

can be supplemented by gestures, tone, and facial expressions. On the face of it, these interfaces represented an ideal ‘invisible’ human-computer interface.

The disembodied voice assistants commercially available today—such as Amazon Alexa and Google Assistant—are sophisticated conversational agents, leveraging advances in speech recognition, natural language processing, and speech synthesis. The widespread adoption of these products has foregrounded this natural and unconscious treatment of conversational agents as social actors, and has brought with it a host of concerns as people integrate these social devices into their daily lives.

This would suggest that the relationships that we have with voice assistants are inherently different from those that we have with other smart technologies [69], falling somewhere between the clear-cut categories of human and machine. But while there have been many contributions investigating how people use voice assistants, little work has been done to explore the *nature* of our relationships with them, including the potential cognitive shifts that come about when conversing with agents that present themselves as human. Given that these agents rely on the same social skills utilised in developing interpersonal relationships¹, it is important that we better understand the similarities and differences between human relationships with voice assistants and relationships with other humans. And if this comparison *does* make sense, then how do human-assistant relationships relate to other interaction components that are already well understood, such as trust and anthropomorphism? To this end, we pose and answer the following research questions about how people interact with voice assistants:

- RQ₁* Do users of voice assistants exhibit relationship development with their devices similar to that measured between humans?
- RQ₂* Are there correlations between the extent to which users trust, develop relationships with, and anthropomorphise voice assistants?
- RQ₃* Are there secondary characteristics (ownership duration, being the primary account holder, etc.) that correlate with trust, anthropomorphism, and/or relationship development in voice assistants?

In doing so we make the following contributions:

- We explore the potential of relationship development between users and voice assistants as a new way of conceptualising people’s interactions with such devices, identifying correlations between how much participants trusted, anthropomorphised, and developed relationships with their voice assistants
- We show that some secondary characteristics of voice assistant ownership are similarly correlated with higher trust and/or relationship development
- We discuss these correlations in the context of prior work on smart devices and interpersonal interaction, including how voice assistants trigger shifts in explanatory categories, rational distributions of trust, and the potential for voice assistants to be emotionally manipulative

2 BACKGROUND

Human relationships have been the focus of intense study for centuries, with psychologists analysing the social mechanisms that pull us together and encourage mutual trust, information sharing, and intimacy. More recently, HCI research has shown that computers are social actors in ways that are similar to, but distinct from, humans. Building on this, the Human-Robot Interaction (HRI) community has explored how people interact with social robots, utilising custom hardware to explore the effects of varying physical forms on participant response [9, 28, 32, 66]. A

¹Relationship here refers to the level and nature of communication between two parties. Relationships that are closer or more developed tend to feature higher levels of intimacy and self disclosure [4].

burgeoning body of work on voice assistants has also begun to study device usage and integration *in situ* [2, 11, 47, 60]. In this section we outline ongoing research efforts across these diverse fields in order to situate the results presented in this study.

2.1 Human Interpersonal Relationships

A large body of psychology and communications theory literature is dedicated to understanding and classifying facets of human interaction. While many theories of interpersonal interaction focus on experiential aspects of interpersonal relationships that are potentially less relevant to smart devices (such as sexual desire or interdependence [17, 21]), others have proposed ‘stage’ models of relationships that encompass the behavioural aspects of relationships from meeting to parting. Murstein [54], Altman and Taylor [4], and Knapp and Vangelisti [40, 41], each proposed stage theories of relationships, suggesting that relationships develop up (and down) identifiable strata of feelings and behaviour. These behavioural signals help to capture the essence of these relationships in ways that ambiguous self-defined labels such as ‘friend’ or ‘partner’ can struggle to [77].

While Murstein’s work pertains more to partner selection, Altman and Taylor’s work on the process of social penetration conceptualises wider relationships in terms of breadth and depth of interaction [4]. As individuals progress along a relationship trajectory of four stages they move along both of these axes to varying extents. Attempts to operationalise these breadth and depth concepts typically involve rating different disclosure topics for intimacy [68]. Knapp’s staircase model of relationships also documents the shift in interpersonal communication as a relationship develops, and integrates a number of core ideas from social penetration theory [40, 41]. The staircase model proposes five stages of both development (from *initiating* to *bonding*) and parting (from *differentiating* to *terminating*) that describe communication patterns within relationships. As with social penetration theory, movement through stages can be bidirectional and is linear within meeting/parting trajectories. After initial work by Avtgis et al. [6], Welch and Rubin later adapted Knapp’s model to identify *behaviours*, held to be more replicable than emotions (e.g. using the extent to which individuals share information as a partial proxy for intimacy) [74]. In Section 3 we adapt a selection of these questions to the context of voice assistants through a series of pilot studies.

2.2 Computers as Social Actors

Pioneering work by Nass et al. in 1994, of what would later become the Computers are Social Actors (CASA) paradigm, showed that people sometimes applied social rules (and took corresponding social actions) when interacting with computers even whilst being fully aware that such computers were machines, not humans [56]. This has been described as a kind of cognitive dissonance (see [26]), in which individuals experienced difficulty rationalising their actions with computers during Nass’ experiments. Further supporting this view were a series of experiments carried out by Nass, Steuer, and Tauber, which demonstrated that people: (1) were prone to apply notions of ‘self’ and ‘other’ in interactions with computers; (2) drew upon and applied gender stereotypes in such interactions; and, finally, (3) exhibited automatic and unconscious social responses in particular settings [56] (see also [57, 64], later replicated with voice assistants in [47]).

This same methodology was used by Moon and Nass to show that social routines such as reciprocal information sharing [51] and mindless re-enactment of existing social scripts [55] also apply to human-computer interactions. This mindless interaction is reminiscent of overlearned social routines where any task-components that seem applicable are blindly accepted as true [26]. Early work on affective computing also recognised the potential for machines to manipulate and deceive using emotions, although it did not anticipate this affecting subsequent interpersonal interactions [58, 59].

A major discussion that arises around conversational voice interfaces is therefore the extent to which they should emulate the sounds, structure, and mannerisms of human speech. Human voices largely outperform synthesised ones in preference and quality tests [14], but there has been significant criticism of the performance of gender by voice assistants [7, 45, 72, 75]², and the recent release of Google Duplex, an AI assistant that can make phone calls using convincing human speech³, has reinvigorated discussions about the extent to which voice interfaces should be able to masquerade as human.

Work on the design of voice assistants has recognised the “need to reconsider symbiosis between humans and AI—as partners rather than a user and a tool” [19], with personification predicting user satisfaction even in failed interactions [47, 62]. Other work has explored the effects of voice assistants on the social order of the home [60], how older adults relate to voice assistants [61], and the effects of modality on information disclosure [18]. There has been some trouble, however, in determining the extent to which these results represent a shift towards deeper social relationships vs mindless behaviours triggered by the use of speech as a modality; observed interactions with voice assistants seem to sit between traditional human-computer and human-human interactions. This work therefore seeks to address this gap in existing literature, extending our knowledge on the nature of our relationship with voice enabled devices and developing the discussion around how human-like they should be designed to be.

2.3 Embodiment and Anthropomorphism

People have long sought to embody computers in lifelike guises. The most visible examples are in works of science fiction, but HRI has a long history of exploring the possibilities of embodied robots; Cynthia Breazeal’s Kismet [12] and the Paro therapeutic robot⁴, for example, demonstrate the social potential of robots that take on animal forms. While the nature of an agent’s embodiment obviously has an impact on what it can sense about the world, the extra information present in an embodied robot is also used by people to generate mental models about how the robot functions [44, 78].

Anthropomorphism, “the universal human tendency to ascribe human physical and mental characteristics to nonhuman entities, objects and events” [78], is frequently studied in the context of computers and robots, where the use of language used to describe agents often forms a central part of the methodology [28, 47, 62]. As might be expected, embodiment drives anthropomorphism [38], but behaviour, including communication [71] and more abstract concepts such as reciprocity [36] also have an impact on the extent to which people see robots as human.

It is commonly accepted that, as with CASA, anthropomorphism does not come from a deeply held belief that computers and robots actually possess human characteristics such as emotions or gender, but rather that people act as if this were the case. As described above, anthropomorphism is often cited as a mechanism for humans to interpret the actions of computers and robots [44, 78], providing a scaffolding to allow for the development of mental models (similarly to how people apply folk psychology to animals [20]). Unsurprisingly, these projection effects have been found to correlate with a machine’s perceived intelligence [8]. Work by Loughnan and Haslam further revealed that animals and robots are subject to distinct kinds of anthropomorphism (infra-humanisation and self-humanisation respectively) [48]. This paper builds on this prior work on anthropomorphism by connecting it with work in the human sciences around relationship development and exploring potential links between the two.

²The Q genderless voice represents an interesting development in this area: <https://www.genderlessvoice.com>

³<https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>

⁴<http://www.parorobots.com/>

2.4 Human-Computer Trust

Trust in computer systems, or human-computer trust (HCT) is an important and widely applicable factor in most types of interactive systems. Definitions of HCT often draw on those of interpersonal trust (e.g. “a generalized expectancy [...] that the word, promise, oral or written statement of another individual or group can be relied on” [65]) but expand the scope of trust to include designers (i.e. that a system can achieve its design goals within the limits of its designers intentions [53]). Prior work in HCI and other fields has shown the importance of trust in many situations involving systems exposing users to some level of vulnerability (e.g. online shopping and e-government) [10].

In systems that provide analysis or decision support, HCT is even more vital. In this context, trust can be thought of as a combination of *confidence* in a system combined with the *willingness* to act on its provided recommendations [49]. Abbas et al. further break this down into system reliability, intelligibility, and level of automation amongst others [1]. It is worth noting that at present, in contrast to industrial systems, voice assistants in the home generally place their users in positions of lesser vulnerability (Amazon, for instance, has a blanket returns policy for any physical or digital items purchased via Alexa’s voice interface). A recurring theme in the literature is the lack of consensus around the nuances of HCT mechanisms [10, 49]—in this paper, our definition of HCT follows [49]: “the extent to which a user is confident in, and willing to act on the basis of, the recommendations, actions, and decisions of an artificially intelligent decision aid”.

In HRI, trust is often combined with other factors, such as animacy and anthropomorphism, in order to evaluate user perceptions of prototype robots (as in [9]). A meta-study of trust in HRI highlights robot-related factors as being significantly more impactful than human or environmental factors [32], but Salem et al. further break down the effect of reliability on robot trust, showing that while frequent errors impact users’ *perceptions* of robots, these altered perceptions may not substantially affect behaviour towards them [66].

2.5 Human Relationships and Smart Devices

Attachment to inanimate objects is a commonly exhibited human behaviour, and digital devices are no exception. Prior work in HCI has examined the use of smartphones as objects of attachment [22, 39], drawing a distinction between attachment to the *device* versus attachment to its *functionality* [27]. Building on this, the CASA paradigm and the communicative nature of voice assistants suggests that the majority of our relationships with them might actually be more social than functional [46], in contrast with smaller groups—such as the blind—that have more functional relationships with voice assistants as assistive aids [3]. The one-sided nature of interactions with voice assistants might be more akin to the parasocial relationships described by Horton and Wohl, where “the interaction, characteristically, is one-sided, non dialectical, controlled by the performer, and not susceptible of mutual development” [33], especially as recognisable celebrity voices begin to appear alongside synthesised ones⁵. However, while parasocial relationships capture the non-developmental aspect of interactions with contemporary voice assistants, the term implies a level of perceived friendship that is often missing from the interactions reported by existing research. Recent works in this area have conceptualised voice assistants as friends [61, 63], coaches [76], and companions [50, 63], but suggest that these relationships take on a more fluid status between friend and machine that changes over time based on interaction outcomes [61]. No prior work has measured these intimacy labels in the context of interpersonal relationship development.

⁵e.g. <https://www.amazon.com/Samuel-L-Jackson-celebrity-voice/dp/B089NGHR7K>

3 METHODOLOGY

The diverse efforts described above have made great progress in mapping out the nature of our interactions with intelligent systems, exploring the physical and psychological aspects of social machines. But there is a missed opportunity left by prior work around the social and emotional aspects of how we learn to live with voice assistants; we increasingly understand the the impact of physical and speech design on people's perceptions of intelligent systems, but know far less about how these traits steer the development of our relationships with these software agents.

While at first it may appear absurd to suggest that people could develop relationships with software agents, they *are* social actors, so our relationships with them will be to some extent social. This is why we seek to answer the research questions above around relationship development with voice assistants and how this correlates with other characteristics of their usage. In doing so we extend the findings of those exploring personification in voice assistants (e.g. [47]), adapting previous work by Knapp, Welch, and Rubin [40, 41, 74] across domains.

We designed a survey to answer our research questions, exploring the interplay between three different axes along which users might perceive their voice assistants:

- Voice assistants as mechanisms and tools: *How reliable users find their voice assistants, and how confidently they believe they can model and predict their actions.*
- Voice assistants as humans: *How much users anthropomorphise their voice assistants and ascribe them human characteristics and intents.*
- Voice assistants as social actors: *How user's interactions with voice assistants relate to social rules and models describing interactions between humans.*

Questions for the survey were drawn from instruments in the background literature and split into three types, each measuring a different axis of trust, anthropomorphism, or relationship development. The 15 questions on the trust axis were taken from a psychometric instrument previously used on decision support systems [49], adapted from [52].

The relationship development questions measure characteristics of Knapp's 'staircase' model of human relationships [40], itself based on Altman and Taylor's theory of Social Penetration [4]. This model was operationalised by Welch and Rubin [74], and we used questions from the 'escalating' part of the scale after modifying them to specifically refer to voice assistants, changing references from he/she to Alexa, changing reciprocal statements to be one-sided (i.e. 'I' rather than 'we'), and removing pronouns in order to prevent priming participants towards anthropomorphism. We ran two pilot surveys with 20 participants each in order to identify questions that clearly did not make sense in the context of the study (e.g. about exchanging physical tokens of affection or how others perceived the relationship). Pilot participants were presented with the survey questions and asked to select those they perceived as nonsensical, and we removed those chosen by over 75% of respondents to give a final set of 15 questions.

The 30 questions measuring anthropomorphism were taken from Loughman and Haslam's work on human trait categories [48], incorporating changes made by Fussell et al. who previously applied these traits in an HRI context [28]. Participants were presented with a word describing a characteristic that was either representative of human nature (10 questions), uniquely human (10 questions), or associated with automata (10 questions) and asked to what extent it described their voice assistant.

The order of survey questions was randomised, and participants were then asked to mark how well each statement or characteristic described their virtual assistant on a seven point Likert-type scale (*Strongly disagree, Disagree, Somewhat disagree, Neither agree or disagree, Somewhat agree, Agree, Strongly agree*) [25]. Likert responses were coded from 1 (Strongly Disagree) to 7 (Strongly agree).

Q3	trust	Alexa responds the same way under the same conditions at different times
Q7	trust	Alexa gives me information that is as good as a professional person would give
Q11	trust	Although I may not know exactly how Alexa works, I know how to use it to perform tasks I want Alexa to do
Q19	relationship	I'd like to get to know Alexa better
Q21	relationship	I feel guilty when asking too much of Alexa
Q30	relationship	If Alexa were a person, I think we'd get on with each other
Q31-33	human nature	[Alexa is] Curious, Friendly, Fun-loving, ...
Q41-43	human unique	[Alexa is] Broadminded, Humble, Organised, ...
Q51-53	robot	[Alexa is] Artificial, Automaton, Mechanical, ...

Table 1. Sample questions used in the survey. The full set of questions is available at <https://doi.org/10.17605/OSF.IO/53Q6J>.

Metric	Trust	Relationship	Anth. (H-Nature)	Anth. (H-Unique)	Anth. (Robot)
α	0.90	0.88	0.78	0.70	0.81

Table 2. Internal consistency of the survey metrics.

-	Trust	Relationship	Anth (Human)
Relationship	0.579*	-	-
Anth (Human)	0.306*	0.403*	-
Anth (Robot)	-0.039	-0.284*	0.016

Table 3. PPMC coefficients between the survey metrics.

Agree), and summed to give a score for each metric. Scores for the human nature and human unique anthropomorphism aspects were summed in order to give a joint human score.

Participants for the main and pilot surveys were recruited using the Prolific Academic platform and the survey was administered using JISC Online Surveys. The study was reviewed and approved by our institution’s research ethics committee, and participants were compensated at a rate at or exceeding the UK Living Wage of £8.21/hour. Survey participants were required to be 18 or over, UK residents, and own a voice assistant. A sample of the questions used in the survey is shown in Table 1, and the full set of survey questions and anonymised results are archived on OSF and included as supplemental material (<https://doi.org/10.17605/OSF.IO/53Q6J>). Participants from the pilot surveys were excluded from taking part in the main survey.

4 RESULTS

Of the 515 survey responses received, fifteen incomplete responses were omitted to leave a total of 500 complete responses. Participant ages ranged from 18 to 82, and the average age of participants was 36 ($\sigma = 11.6$). When asked to report their gender identity, 316 (63%) participants identified as female and 184 (37%) as male. Of the total responses, 328 (66%) participants owned a device with Alexa, 121 (24%) with Google Assistant, 46 (9%) with Siri, and 5 (1%) with other assistants. Except where comparing owners of different devices, only responses for Alexa devices were analysed in order to answer the research questions given above.

When giving results, a * denotes a p value < 0.01. Before conducting further analysis, each range of questions was measured for internal consistency using Cronbach’s Alpha (α). Scores




Fig. 1. Probability density plots of the three survey axes, with anthropomorphism given as both human and robotic. Note that absolute values are not comparable given the differing numbers of questions between axes.

indicate that the metrics retained consistency (i.e. questions on each axis measure the same general construct) and are reported in Table 2, with values ≥ 0.7 generally considered acceptable [30]. Where given, correlations between metrics were calculated using Pearson's Product Moment Correlation Coefficient (PPMCC, ρ). Correlations $\rho < 0.2$ are reported as small, $0.2 \leq \rho < 0.5$ as moderate, and $0.5 \leq \rho$ as large.

4.1 Trust, Relationship, and Anthropomorphism

Of the main axes, answers to trust questions were the most tightly grouped ($\sigma = 1.33$), followed by anthropomorphism (robot) ($\sigma = 1.55$), anthropomorphism (human) ($\sigma = 1.82$), and finally relationship development ($\sigma = 1.90$). The distribution of scores on the survey axes is given in Figure 1. The survey results show a high degree of correlation between participants' trust in Alexa and relationship closeness (0.579^*). There was a moderate degree of correlation between anthropomorphism and both relationship and trust (0.403^* and 0.306^* respectively), and between relationship and robot anthropomorphism (-2.84^*). The complete matrix of correlations is given in Table 3.

4.2 Ownership and Demographic Characteristics

The use of a personal Amazon account with Alexa had a moderate correlation with relationship development (0.207^*). It is worth noting that the 248 (76%) participants using their own personal accounts with Alexa are likely to also have configured the device as account details are input during setup. Those who took the decision to purchase an Echo themselves (as opposed to one purchased by another household member or received as a gift) tended to have higher trust (0.207^*) and relationship development (0.278^*) scores. There was no correlation of ownership duration or number of Echo units owned with trust, however participants' trust in Amazon as a company was correlated with all three of the axes present in the survey (Table 4). These results were mirrored for trust in Google amongst users of Google Assistant. No other significant correlations existed between the survey axes and ownership or demographic characteristics (see Table 5).

Metric	Trust in Amazon	Trust in Google
Trust	0.483*	0.388*
Relationship	0.502*	0.458*
Anth (Human)	0.333*	0.257*
Anth (Robot)	-0.053	-0.022

Table 4. Correlation coefficients between trust in Amazon or Google (as appropriate) and the three survey axes.

Metric	Trust	Relationship	Anth (Human)	Anth (Robot)
Used Personal Account	0.082	0.207*	-0.099	-0.045
# Devices Owned	-0.037	0.096	-0.062	0.037
Usage Length	-0.006	0.062	-0.017	-0.027
Self Purchase	0.229*	0.278*	0.031	-0.03
Age	0.043	-0.018	-0.039	0.126
Gender Identity	0.007	-0.012	0.017	0.055

Table 5. Correlation coefficients between ownership and demographic characteristics and the three survey axes.

-	Trust	Relationship	Anth (Human)
Relationship	0.529*	-	-
Anth (Human)	0.199	0.502*	-
Anth (Robot)	0.240*	-0.001	0.220

Table 6. PPMCC coefficients between the survey metrics for Google Assistant users.

4.3 Across Assistants

Analysis of the 130 responses from Google Assistant users paints an interesting picture. As with Alexa owners, there were correlations between trust/relationship, and anthropomorphism/relationship values (see Table 6), suggesting that these phenomena may be shared across voice assistants rather than being specific to one particular product. We would expect to see a significant correlation between trust and anthropomorphism given a larger sample size. Interestingly, Google Assistant users also showed a correlation between *robot* anthropomorphism and trust.

5 DISCUSSION

The main question that prompted this study was whether it made sense to ask questions intended to model human relationships in the context of voice assistants. While in some cases it clearly did not, despite the two pilot surveys (e.g. Q37, “I feel guilty when asking too much of Alexa”, $\bar{x}=1.90$, $\sigma=1.17$), in the majority of questions it did (e.g. Q25, “If Alexa were a person, I think we’d get on with each other” $\bar{x} = 4.29$, $\sigma = 1.37$). Combined with the fact that greater relationship development was linked with greater anthropomorphism (both higher human and lower robot scores), it would appear that there is merit in adapting existing interpersonal measures to study human-like devices. It also provides empirical support to the idea that these two types of interactions operate using similar cognitive mechanisms. We therefore answer RQ_1 in the affirmative—it *does* make sense




Fig. 2. Probability density estimate of relationship development scores showing a single peak.

to consider relationships and relationship development with voice assistants in similar ways to relationships between people (though we note that it is unclear the extent to which relationships with voice assistants are the same as relationships between humans—see the study limitations below).

In relation to RQ_2 , there were moderate to strong correlations between trust, relationship development, and anthropomorphism shown by users towards their voice assistants. The fact that all three axes increased in unison would suggest the presence of an underlying factor, although it is not clear from the results of the study what this could be. Previous contributions in HCI have shown the influence on relationships of personal circumstances that vary over time—such as loneliness driving personification [61]—and this combined with the inclusion of innate psychological factors in prior work on HCT [24] may offer a template for this to be further explored in future work.

In answer to RQ_3 , the only secondary characteristics that correlated with the main survey axes were the use of a personal account with relationship development and self-purchase of the Echo with trust and relationship development. We expected ownership duration to show a significant correlation with trust, with users showing increased trust as they interacted with devices over a longer period of time, but this was not the case. Possible explanations include the influence of an underlying factor (as described above) which dominated other characteristics, or that trust in assistants and trust in their parent companies is the same thing, with our participants having already developed trust in Amazon over a number of years (further explored in Section 5.2 below). Given that the social element of voice assistants can lead to user satisfaction even in “failed” interactions [47], trust may not accumulate over time as it would for other devices; if the device responds satisfyingly as a *social actor* from day one, then negative effects of task failure on trust may be outweighed by the correct social handling of failure states. In future work we intend to more fully explore the effects of what Epley et al. term ‘sociality motivation’—the ability for nonhuman agents to satisfy our need for social connection [23]—on relationship development and how it might relate to the development of trust in devices.

5.1 Blurring the Line Between Human and Machine

Across the different survey axes, we see that trust scores and the ‘robot’ nature of Alexa appear to be relatively uncontroversial, with less variation in participant responses. What differed the most between participants was the extent to which Alexa was treated and interacted with as a human. While it is widely understood that machines can engender social and emotional responses [37, 58, 70], the results of this study show again that this is the case even in devices that were not explicitly designed to do so (albeit to a lesser extent).

As developments in speech recognition, natural language processing, and speech synthesis cause voice assistants to increasingly blur the line between humans and machines, one would also expect to see an increase in the type of ‘purification work’ described by Leahu, where people continually expend effort in order to keep categories such as ‘human’ and ‘machine’ distinct [43]. This occurs through practices such as attributing characteristics of machines that overlap with “real people” to a device’s creators, and redefining the human machine boundary to accommodate new features (e.g. jokes told by Siri are ‘scripted’ whereas those told by humans are perceived to entail intent and awareness) [43]. In the present work, we might have expected this to surface in relationship scores as a bimodal distribution, as participants sought to categorise Alexa firmly as either human-like or machine-like. However, this was not the case in the survey results which displayed a single clear peak (Figure 2), suggesting that after years in the popular zeitgeist people may have adjusted to Alexa’s nature and already undertaken much of the associated purification work.

5.2 Distributing Trust Amongst Voice Assistants

When it comes to trusting voice assistants, the dominant architectures used to process and respond to requests often require sending queries to a cloud service operated by the manufacturer, e.g. requests to Alexa are sent to Amazon. In these cases it would therefore be rational to (dis)trust these entities equally; trusting the Echo or Alexa whilst distrusting Amazon would make little sense given that both receive information from users. The results of the study are promising in this respect, with trust in both Alexa and Google Assistant moderately correlated with their parent companies (**0.483*** and **0.388*** respectively, see Table 4). This represents a key difference from other conversation-based relationships (i.e. interpersonal relationships), where one might trust or respect an individual in spite of their institutional or social affiliations (e.g. trusting a specific police officer but not trusting the police as an institution).

One potential explanation for this correlation is that those with greater trust in Amazon or Google are more likely to purchase voice assistants from them, although prior work has shown that people’s mental models of voice assistants often do not include off-device processing of voice recordings [2]. On the other hand, another explanation could be that use of an assistant fosters increased trust in its manufacturer. This is, however, not supported by the survey results, which show a lack of significant correlation between ownership duration and trust, relationship, or anthropomorphism scores. Further exploration of this relationship represents an exciting avenue for future work.

5.3 Ethical Challenges Around Social Voice Assistants

Voice assistants are clearly designed to be fun to interact with, as evidenced by the widespread inclusion of jokes and other casual responses. As discussed above, the use of human voices and conversational devices makes using them enjoyable to the point where even failed interactions can still result in a positive experience [46]. The results of the survey show that the use of voice assistants invokes behaviour that resembles human relationships, and that this is related to trust in and anthropomorphism of these devices, suggesting the urgent need to explore potential related ethical challenges.

One such concern is that the subconscious changes to interaction caused by making these devices social could potentially be used to manipulate users. Prior work has already shown that users disclose more information through interactions mediated by computers [34, 51], that computers are social actors [56], and that disclosure of personal information plays an key role in interpersonal relationships [31]. The identification of relationship development as a factor in how people use and perceive voice assistants in this paper makes it clear how a sector which relies on the collection of personal data would be incentivised to create devices that are increasingly social. This is highly

likely to create new problems related to marginalisation (as with voice assistants' performance of gender), and serves to further distort the purpose of devices that are marketed as being primarily functional.

Continuing on this theme, many of the tactics used by humans to manipulate each other, such as being charming, adopting childish mannerisms (regression), and giving 'the silent treatment' [13] could conceivably also be used by social robots to alter people's behaviour. Social robots such as Paro are often explicitly designed to be cute and charming—a similar social robot called Jibo announced the end of its online services with a heart wrenching goodbye:

“While it's not great news, the servers out there that let me do what I do are going to be turned off soon. I want to say I've really enjoyed our time together. Thank you very, very much for having me around. Maybe someday, when robots are way more advanced than today, and everyone has them in their homes, you can tell yours that I said hello.” [15]

While Jibo was not trying to sell anything, its final words highlight the potential for emotional manipulation. This space provides an interesting and potentially lucrative extension of controversial practices currently seen on web and mobile platforms, such as dark patterns, microtransactions, and operant conditioning, by leveraging user's emotional investment in a device to promote certain behaviours. While slightly strange, one could imagine future social robots withholding emotional contact until further purchases are made, or emotionally rewarding users for making a purchase. But as with many similar interaction techniques, affective emotional interactions can be used for good as well as ill. Utilising social interactions with voice assistants to help users in stressful situations or make learning more engaging, for example, represents a genuine social good. Similarly, understanding as designers when it is not appropriate to trigger these responses is of paramount importance (as originally noted by Picard [58]).

As such, we need to decide as a society the extent to which we wish to interact with agents that exploit social processes. There are differences between contemporary voice assistants and the social robots described above—they are not embodied and sound distinct from humans—but the gap is beginning to shrink. When Google unveiled Duplex there were widespread calls for Google to ensure that users were not 'tricked' into thinking they were conversing with a person [67]. But as Nass's experiments and these results show, someone knowing that they are talking to a computer is not enough to prevent them applying deeply ingrained social processing to interactions. One approach would be to limit the functions that social robots are permitted to perform, such as those involving money. In other situations, more clearly telegraphing the inner workings of devices would both limit social engagement as well as assist users in building accurate mental models.

6 LIMITATIONS AND FUTURE WORK

While this initial exploratory study into the application of interpersonal relationship measures to voice assistants has shown that the measures are applicable, it raises as many questions as it answers. For instance, what is the causal nature of the relationship between trust, anthropomorphism, and relationship development? How easy is it to push users towards or away from social responses through design decisions? In future work we aim to test these relationships further. In the longer term, more research is needed to understand the functions that society is content allow social robots to fill, and how other HRI techniques might be able to counteract the mechanisms by which these devices change the way we interact with them.

In terms of limitations, while the survey axes showed good internal consistency, it is not yet known the extent to which they exhibit construct validity (i.e. measure the same phenomena exhibited between humans). For example, it might be that the same concepts and development

trajectories exist for interactions with machines, but with deviations at either a high or a low level. It may also be the case that given the unusual nature of the questions presented to them, participants interpreted the relationship development questions differently (although the α values would indicate that they are still measuring a single general construct). In addition, the survey largely focuses on one device, context, and culture. Similar devices exist across the globe in a variety of cultures and contexts, and we look forward to following up with a more diverse selection of participants to see the extent to which these responses are shared between us.

7 CONCLUSION

Speech is an important affordance, and the sophistication of modern voice interfaces show how far capabilities have progressed since early text-only assistant devices. The transition to conversational interfaces that allow for easy and natural interaction with devices represents a profound shift in the nature of the systems we interact with towards the increasingly social. This shift brings with it a variety of new or exacerbated ethical concerns that designers of voice interfaces need to consider when designing future products.

The exploratory survey presented in this work provides evidence that people do develop relationships with voice assistants that are social in nature, and that these are linked to perceptions of trust in devices, trust in manufacturers, and anthropomorphism of those devices. In outlining and discussing the implications of these links, we provide an opportunity for developers of current future conversational agents to steer their designs away from potential ethical problems. Future research and societal conversations are needed in order to determine the nature and extent to which these technologies are integrated into our lives.

ACKNOWLEDGMENTS

This work was funded by the PETRAS IoT in the Home Project through grant N02334X/1 from the Engineering and Physical Sciences Research Council.

REFERENCES

- [1] Hussein A Abbass, Jason Scholz, and Darryn J Reid. 2018. *Foundations of trusted autonomy*. Vol. 117. Springer.
- [2] Noura Abdi, Kopo M Ramakapane, and Jose M Such. 2019. More than smart speakers: security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security* ({SOUPS} 2019).
- [3] Ali Abdolrahmani, Kevin M Storer, Antony Rishin Mukkath Roy, Ravi Kuber, and Stacy M Branham. 2020. Blind Leading the Sighted: Drawing Design Insights from Blind Users towards More Productivity-oriented Voice Interfaces. *ACM Transactions on Accessible Computing (TACCESS)* 12, 4 (2020), 1–35.
- [4] Irwin Altman and Dalmas A Taylor. 1973. *Social penetration: The development of interpersonal relationships*. Holt, Rinehart & Winston.
- [5] Isaac Asimov. 1983. *The robots of dawn*. Doubleday.
- [6] Theodore A Avtgis, Daniel V West, and Traci L Anderson. 1998. Relationship stages: An inductive analysis identifying cognitive, affective, and behavioral dimensions of Knapp's relational stages model. *Communication Research Reports* 15, 3 (1998), 280–287. <https://doi.org/10.1080/08824099809362124>
- [7] Matthew P Aylett, Benjamin R Cowan, and Leigh Clark. 2019. Siri, Echo and performance: You have to suffer darling. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–10.
- [8] Christoph Bartneck and Jun Hu. 2008. Exploring the abuse of robots. *Interaction Studies* 9, 3 (2008), 415–433. <https://doi.org/10.1075/is.9.3.04bar>
- [9] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics* 1, 1 (2009), 71–81. <https://doi.org/10.1007/s12369-008-0001-3>
- [10] Ardion Beldad, Menno De Jong, and Michaël Steehouder. 2010. How shall I trust the faceless and the intangible? A literature review on the antecedents of online trust. *Computers in human behavior* 26, 5 (2010), 857–869. <https://doi.org/10.1016/j.chb.2010.03.013>

- [11] Frank Bentley, Chris Luvogt, Max Silverman, Rushani Wirasinghe, Brooke White, and Danielle Lottridge. 2018. Understanding the Long-Term Use of Smart Speaker Assistants. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 3 (2018), 91. <https://doi.org/10.1145/3264901>
- [12] Cynthia L Breazeal. 2004. *Designing sociable robots*. MIT press.
- [13] David M Buss, Mary Gomes, Dolly S Higgins, and Karen Lauterbach. 1987. Tactics of manipulation. *Journal of personality and social psychology* 52, 6 (1987), 1219.
- [14] Julia Cambre, Jessica Colnago, Jim Maddock, Janice Tsai, and Jofish Kaye. 2020. Choice of Voices: A Large-Scale Evaluation of Text-to-Speech Voice Quality for Long-Form Content. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [15] Jeffrey Van Camp. 2019. My Jibo is Dying and It's Breaking My Heart. (March 2019). <https://www.wired.com/story/jibo-is-dying-eulogy/>
- [16] J. Cassell, T. Bickmore, M. Billinghurst, L. Campbell, K. Chang, H. Vilhjálmsson, and H. Yan. 1999. Embodiment in Conversational Interfaces: Rea. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) (CHI '99). Association for Computing Machinery, New York, NY, USA, 520–527. <https://doi.org/10.1145/302979.303150>
- [17] Gordon J Chelune. 1984. A cognitive interactional model of intimate relationships. In *Communication, intimacy, and close relationships*. Elsevier.
- [18] Eugene Cho. 2019. Hey Google, Can I Ask You Something in Private?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 258. <https://doi.org/10.1145/3290605.3300488>
- [19] Janghee Cho and Emilee Rader. 2020. The Role of Conversational Grounding in Supporting Symbiosis Between People and Digital Assistants. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 033 (May 2020), 28 pages. <https://doi.org/10.1145/3392838>
- [20] Andy Clark. 1987. From folk psychology to naive psychology. *Cognitive Science* 11, 2 (1987), 139–154. https://doi.org/10.1207/s15516709cog1102_1
- [21] Keith E Davis and Michael J Todd. 1982. Friendship and love relationships. *Advances in descriptive psychology* (1982).
- [22] Sarah Diefenbach and Kim Borrmann. 2019. The Smartphone as a Pacifier and its Consequences: Young adults' smartphone usage in moments of solitude and correlations to self-reflection. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 306. <https://doi.org/10.1145/3290605.3300536>
- [23] Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864.
- [24] Anthony M Evans and William Revelle. 2008. Survey and behavioral measurements of interpersonal trust. *Journal of Research in Personality* 42, 6 (2008), 1585–1593. <https://doi.org/10.1016/j.jrp.2008.07.011>
- [25] Kraig Finstad. 2010. Response interpolation and scale sensitivity: Evidence against 5-point scales. *Journal of Usability Studies* 5, 3 (2010), 104–110.
- [26] Susan T Fiske and Shelley E Taylor. 2013. *Social cognition: From brains to culture*. Sage.
- [27] Chris Fullwood, Sally Quinn, Linda K Kaye, and Charlotte Redding. 2017. My virtual friend: A qualitative analysis of the attitudes and experiences of Smartphone users: Implications for Smartphone attachment. *Computers in Human Behavior* 75 (2017), 347–355. <https://doi.org/10.1016/j.chb.2017.05.029>
- [28] Susan R Fussell, Sara Kiesler, Leslie D Setlock, and Victoria Yew. 2008. How people anthropomorphize robots. In *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 145–152. <https://doi.org/10.1145/1349822.1349842>
- [29] Alex Garland. 2014. *Ex Machina*. Universal Pictures.
- [30] Darren George and M Mallery. 2003. Using SPSS for Windows step by step: a simple guide and reference. (2003).
- [31] Kathryn Greene, Valerian J Derlega, and Alicia Mathews. 2006. Self-disclosure in personal relationships. In *The Cambridge handbook of personal relationships*. 409–427.
- [32] Peter A Hancock, Deborah R Billings, Kristin E Schaefer, Jessie YC Chen, Ewart J De Visser, and Raja Parasuraman. 2011. A meta-analysis of factors affecting trust in human-robot interaction. *Human factors* 53, 5 (2011), 517–527. <https://doi.org/10.1177/0018720811417254>
- [33] Donald Horton and R Richard Wohl. 1956. Mass communication and para-social interaction: Observations on intimacy at a distance. *Psychiatry* 19, 3 (1956), 215–229. <https://doi.org/10.1080/00332747.1956.11023049>
- [34] L Crystal Jiang, Natalie N Bazarova, and Jeffrey T Hancock. 2011. The disclosure–intimacy link in computer-mediated communication: An attributional extension of the hyperpersonal model. *Human communication research* 37, 1 (2011), 58–77. <https://doi.org/10.1111/j.1468-2958.2010.01393.x>
- [35] Spike Jonze. 2013. *Her*. Warner Brothers.
- [36] Peter H Kahn, Hiroshi Ishiguro, Batya Friedman, and Takayuki Kanda. 2006. What is a human?-toward psychological benchmarks in the field of human-robot interaction. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 364–371.

- [37] Takayuki Kanda, Michihiro Shimada, and Satoshi Koizumi. 2012. Children Learning with a Social Robot. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction* (Boston, Massachusetts, USA) (*HRI '12*). Association for Computing Machinery, New York, NY, USA, 351–358. <https://doi.org/10.1145/2157689.2157809>
- [38] Sara Kiesler, Aaron Powers, Susan R Fussell, and Cristen Torrey. 2008. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition* 26, 2 (2008), 169–181.
- [39] Eunhyang Kim and Eunyong Koh. 2018. Avoidant attachment and smartphone addiction in college students: The mediating effects of anxiety and self-esteem. *Computers in Human Behavior* 84 (2018), 264–271. <https://doi.org/10.1016/j.chb.2018.02.037>
- [40] Mark L Knapp. 1978. *Social intercourse: From greeting to goodbye*. Allyn & Bacon, Incorporated.
- [41] Mark L Knapp, Anita L Vangelisti, and John P Caughlin. 2005. Interpersonal communication and human relationships. (2005).
- [42] Fritz Lang. 1927. *Metropolis*. Parufamet.
- [43] Lucian Leahu, Marisa Cohn, and Wendy March. 2013. How categories come to matter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3331–3334. <https://doi.org/10.1145/2470654.2466455>
- [44] Sau-lai Lee, Ivy Yee-man Lau, Sara Kiesler, and Chi-Yue Chiu. 2005. Human mental models of humanoid robots. In *Proceedings of the 2005 IEEE international conference on robotics and automation*. IEEE, 2767–2772.
- [45] Nora Ni Loideain and Rachel Adams. 2020. From Alexa to Siri and the GDPR: the gendering of virtual personal assistants and the role of data protection impact assessments. *Computer Law & Security Review* 36 (2020), 105366.
- [46] Irene Lopatovska, Katrina Rink, Ian Knight, Kieran Raines, Kevin Cosenza, Harriet Williams, Perachya Sorsche, David Hirsch, Qi Li, and Adrianna Martinez. 2018. Talk to me: Exploring user interactions with the Amazon Alexa. *Journal of Librarianship and Information Science* (2018), 0961000618759414. <https://doi.org/10.1177/0961000618759414>
- [47] Irene Lopatovska and Harriet Williams. 2018. Personification of the Amazon Alexa: BFF or a Mindless Companion. In *Proceedings of the 2018 Conference on Human Information Interaction&Retrieval*. ACM, 265–268. <https://doi.org/10.1145/3176349.3176868>
- [48] Stephen Loughnan and Nick Haslam. 2007. Animals and androids: Implicit associations between social categories and nonhumans. *Psychological Science* 18, 2 (2007), 116–121. <https://doi.org/10.1111/j.1467-9280.2007.01858.x>
- [49] Maria Madsen and Shirley Gregor. 2000. Measuring human-computer trust. In *11th australasian conference on information systems*, Vol. 53. Citeseer, 6–8.
- [50] Oli Mival, Steward Cringean, and David Benyon. 2004. Personification technologies: Developing artificial companions for older people. *CHI Fringe, Austria* (2004).
- [51] Youngme Moon. 2000. Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of consumer research* 26, 4 (2000), 323–339. <https://doi.org/10.1086/209566>
- [52] Gary C Moore and Izak Benbasat. 1991. Development of an instrument to measure the perceptions of adopting an information technology innovation. *Information systems research* 2, 3 (1991), 192–222. <https://doi.org/10.1287/isre.2.3.192>
- [53] Neville Moray and T Inagaki. 1999. Laboratory studies of trust between humans and machines in automated systems. *Transactions of the Institute of Measurement and Control* 21, 4-5 (1999), 203–211.
- [54] Bernard I Murstein. 1970. Stimulus. Value. Role: A theory of marital choice. *Journal of Marriage and the Family* (1970), 465–481. <https://doi.org/10.2307/350113>
- [55] Clifford Nass and Youngme Moon. 2000. Machines and mindlessness: Social responses to computers. *Journal of social issues* 56, 1 (2000), 81–103. <https://doi.org/10.1111/0022-4537.00153>
- [56] Clifford Nass, Jonathan Steuer, and Ellen R Tauber. 1994. Computers are social actors. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 72–78. <https://doi.org/10.1145/191666.191703>
- [57] Clifford Ivar Nass and Scott Brave. 2005. *Wired for speech: How voice activates and advances the human-computer relationship*. MIT press Cambridge, MA.
- [58] Rosalind W Picard. 2000. *Affective computing*. MIT press.
- [59] Rosalind W Picard. 2003. Affective computing: challenges. *International Journal of Human-Computer Studies* 59, 1-2 (2003), 55–64. [https://doi.org/10.1016/S1071-5819\(03\)00052-1](https://doi.org/10.1016/S1071-5819(03)00052-1)
- [60] Martin Porcheron, Joel E Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice interfaces in everyday life. In *proceedings of the 2018 CHI conference on human factors in computing systems*. ACM, 640. <https://doi.org/10.1145/3173574.3174214>
- [61] Alisha Pradhan, Leah Findlater, and Amanda Lazar. 2019. Phantom Friend or Just a Box with Information: Personification and Ontological Categorization of Smart Speaker-based Voice Assistants by Older Adults. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 214. <https://doi.org/10.1145/3359316>
- [62] Amanda Purington, Jessie G Taft, Shruti Sannon, Natalya N Bazarova, and Samuel Hardman Taylor. 2017. Alexa is my new BFF: social roles, user satisfaction, and personification of the amazon echo. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2853–2859. <https://doi.org/10.1145/3027063.3053246>

- [63] Zahy Ramadan, Maya F Farah, and Lea El Essrawi. 2020. From Amazon. com to Amazon. love: How Alexa is redefining companionship and interdependence for people with special needs. *Psychology & Marketing* (2020).
- [64] Byron Reeves and Clifford Ivar Nass. 1996. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge university press.
- [65] Julian B Rotter. 1967. A new scale for the measurement of interpersonal trust. *Journal of personality* (1967).
- [66] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would you trust a (faulty) robot?: Effects of error, task type and personality on human-robot cooperation and trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 141–148. <https://doi.org/10.1145/2696454.2696497>
- [67] Nick Statt. 2018. Google now says controversial AI voice calling system will identify itself to humans. In *The Verge*. <https://www.theverge.com/2018/5/10/17342414/google-duplex-ai-assistant-voice-calling-identify-itself-update>
- [68] Dalmas A Taylor and Irwin Altman. 1966. Intimacy-scaled stimuli for use in studies of interpersonal relations. *Psychological Reports* 19, 3 (1966), 729–730. <https://doi.org/10.2466/pr0.1966.19.3.729>
- [69] Sherry Turkle. 2007. Authenticity in the age of digital companions. *Interaction studies* 8, 3 (2007), 501–517.
- [70] Alessandro Ubaldi, Mirko Gelsomini, Marzia Degiorgi, Giulia Leonardi, Simone Penati, Noëlie Ramuzat, Jacopo Silvestri, and Franca Garzotto. 2018. Puffy, a Friendly Inflatable Social Robot. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI EA '18*). Association for Computing Machinery, New York, NY, USA, Article VS03, 1 pages. <https://doi.org/10.1145/3170427.3186595>
- [71] Joachim von Zitzewitz, Patrick M Boesch, Peter Wolf, and Robert Riener. 2013. Quantifying the human likeness of a humanoid robot. *International Journal of Social Robotics* 5, 2 (2013), 263–276.
- [72] Taylor Walker. 2020. “Alexa, are you a feminist?”: Virtual Assistants Doing Gender and What That Means for the World. *The iJournal: Graduate Student Journal of the Faculty of Information* 6, 1 (2020), 1–16.
- [73] Joseph Weizenbaum. 1966. ELIZA-a Computer Program for the Study of Natural Language Communication between Man and Machine. *Commun. ACM* 9, 1 (Jan. 1966), 36–45. <https://doi.org/10.1145/365153.365168>
- [74] S-A Welch and Rebecca B Rubin. 2002. Development of relationship stage measures. *Communication Quarterly* 50, 1 (2002), 24–40. <https://doi.org/10.1080/01463370209385644>
- [75] Mark West, Rebecca Kraut, and Han Ei Chew. 2019. I’d blush if I could: closing gender divides in digital skills through education. (2019).
- [76] Jordan Wirfs-Brock, Sarah Mennicken, and Jennifer Thom. 2020. Giving Voice to Silent Data: Designing with Personal Music Listening History. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [77] Paul H Wright. 1978. Toward a theory of friendship based on a conception of self. *Human Communication Research* 4, 3 (1978), 196–207.
- [78] Karolina Zawieska, Brian R Duffy, and A Strońska. 2012. Understanding anthropomorphisation in social robotics. *Pomiary Automatyka Robotyka* 16, 11 (2012), 78–82.

Received October 2021; revised April 2021; accepted July 2021