



Do You Mind? User Perceptions of Machine Consciousness

Downloaded from: <https://research.chalmers.se>, 2025-12-05 01:47 UTC

Citation for the original published paper (version of record):

Scott, A., Neumann, D., Niess, J. et al (2023). Do You Mind? User Perceptions of Machine Consciousness. Conference on Human Factors in Computing Systems - Proceedings.
<http://dx.doi.org/10.1145/3544548.3581296>

N.B. When citing this work, cite the original published paper.

Do You Mind? User Perceptions of Machine Consciousness

Ava Scott

ava.scott.20@ucl.ac.uk

University College London

UCL Interaction Centre (UCLIC) and Institute of Cognitive
Neuroscience (ICN)
London, United Kingdom

Jasmin Niess

jasminni@uio.no

University of Oslo

Department of Informatics

Oslo, Norway

University of St. Gallen

Institute of Computer Science
Switzerland

Daniel Neumann

daniel.neumann@unisg.ch

University of St. Gallen

Institute of Computer Science

St. Gallen, Switzerland

Ludwig Maximilian University of Munich

Institut für Informatik

Munich, Germany

Paweł W. Woźniak

pawel.wozniak@chalmers.se

Chalmers University of Technology

Department of Computer Science and Engineering

Gothenburg, Sweden




Figure 1: These two cartoons depict interactive dilemmas posed by the existence of machines that appear conscious. Artwork by Julia Minorowicz

ABSTRACT

The prospect of machine consciousness cultivates controversy across media, academia, and industry. Assessing whether non-experts perceive technologies as conscious, and exploring the consequences of this perception, are yet unaddressed challenges in Human Computer Interaction (HCI). To address them, we surveyed 100 people, exploring their conceptualisations of consciousness and if and how they perceive consciousness in currently available interactive technologies. We show that many people already perceive a degree of consciousness in GPT-3, a voice chat bot, and a robot vacuum cleaner. Within participant responses we identified



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

CHI '23, April 23–28, 2023, Hamburg, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9421-5/23/04.

<https://doi.org/10.1145/3544548.3581296>

dynamic tensions between denial and speculation, thinking and feeling, interaction and experience, control and independence, and rigidity and spontaneity. These tensions can inform future research into perceptions of machine consciousness and the challenges it represents for HCI. With both empirical and theoretical contributions, this paper emphasises the importance of HCI in an era of machine consciousness, real, perceived or denied.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models.**

KEYWORDS

Consciousness, Machine Consciousness, Technology Consciousness

ACM Reference Format:

Ava Scott, Daniel Neumann, Jasmin Niess, and Pawel W. Woźniak. 2023. Do You Mind? User Perceptions of Machine Consciousness. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*, April 23–28, 2023, Hamburg, Germany. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3544548.3581296>

1 INTRODUCTION

Digital technologies are increasingly able to emulate the perceptive and behavioural repertoire of living things, including humans. High-profile stories of such technologies, both fictional [2, 116] and real [124], have supercharged the controversy surrounding the attribution of consciousness to computers [82, 120]. In this paper we argue that, rather than focusing on the existence of machine consciousness, scholars of human computer interaction (HCI) should focus on peoples' perception of machine consciousness (PMC). The complex ontological problem of consciousness is not within the remit of HCI research and is instead addressed by neuroscientists and philosophers [26, 64, 106]. In the absence of direct evidence or clarity, users develop mental models of how a technology works, and how this implicates their interaction [20]. One such model could result in the perception of consciousness. Thus, from a HCI perspective, the first step is not to confirm nor deny the existence of machine consciousness. Instead, it is crucial to understand if and how machine consciousness is perceived by users, how this may affect the interaction between humans and computers, and how we can mitigate the consequences of such perception for users and society. As a discipline with extensive skills in measuring and understanding user experience, an inherently subjective phenomena, HCI has the relevant expertise and methodology to pursue this agenda [70].

What could it mean to perceive consciousness in a machine? While modern neuroscience seems to be closing in on theoretical and empirical concepts of 'consciousness' [64, 106], the term continues to court controversy and conflation. Consciousness is often understood as the phenomenon of subjectivity, of 'having experience' and, at least for humans, also being aware of that experience. However, there is a long list of concepts that are adjacent or equivalent to that of consciousness. Those that have been investigated within HCI research include the concepts of sentience [79, 117], awareness [23, 61], intentionality [101], machine emotion [35, 59, 130],

and theory of mind [49, 60]. While these concepts have been identified, there has been no structured, theoretical or empirical understanding of how people may conceptualise, perceive or deny machine consciousness in currently available devices and systems. This paper hopes to begin this line of inquiry.

Machine consciousness, if perceptible in an interface, can have significant consequences for the interaction between humans and technologies. For example, a software engineer was fired from Google last year, after he perceived the LaMDA chatbot as self-aware and declared this to the world [124]. In everyday life, some users interpret their vacuum robot cleaners as 'enlightened' [45], and debate whether home assistants are deserving of respect or empathy [30, 92]. However, there has not yet been an explicit investigation into the interactional consequences of PMC. This perception could lead to rich, complex relationships of empathy, including companionship, consideration, and concern [37, 87], as well as of hostility, with fears of competition and domination [11, 78]. As well as demonstrating PMC as a distinct phenomena, this paper makes an initial investigation into these contrasting narratives, and the tensions between them.

This paper reports on non-experts' understandings of machine consciousness in currently available interactive systems, and assesses the term's colloquial associations. First, we conducted a literature review on the term 'machine consciousness' and related concepts in HCI. Then, in a 15-question survey of $n = 100$ participants, we explore two key research questions. Firstly, do people perceive machine consciousness in contemporary technologies? If so, to what degree? Secondly, what are the narratives, concerns, and speculations that people currently associate with machine consciousness? To make these questions more tangible and relevant for our participants, we showed them three short videos, introducing Amazon's Alexa, OpenAI's GPT-3, and a vacuum robot, respectively [65]. These research probes were chosen for their diversity in features and familiarity, as well as their presence in the media [18, 31, 45] and in literature concerning consciousness [39, 64, 83, 87, 106, 110].

Using these findings, we extracted five dynamic tensions within PMC that illustrate conceptual, pragmatic, and philosophical complexities. The dynamic tensions are as follows: denial and speculation, thinking and feeling, interaction versus experience, being controlled versus independent, and rigidity versus spontaneity. We use these tensions to inform four challenges for HCI in an era of perceived machine consciousness. Widely varying positions of scepticism, indifference, and curiosity towards machine consciousness may fuel negotiations of hierarchy, equality and control between humans and technologies. Knowledge of these positions will help interaction designers better navigate the opportunities and risks of machine consciousness, whether that be real, perceived or denied.

This paper contributes the following: i) theoretical motivation for HCI to explore *perceived* machine consciousness (PMC) rather than machine consciousness outright, ii) empirical evidence for the perception of machine consciousness ($n = 100$), iii) the identification of dynamic tensions that could inform future research and design into the phenomena of PMC, and, iv) future challenges for HCI in an era of PMC. Overall, this paper confirms that some people already perceive consciousness in machines, and presents a theoretical launchpad for future HCI research into PMC.

2 BACKGROUND & RELATED WORK

We first introduce the HCI reader to concepts and theories of consciousness which are prevalent in recent literature outside of HCI. We then turn to research on perceptions of consciousness, primarily from the psychological sciences. Then, we review the use and study of concepts related to consciousness in HCI.

2.1 Neuroscientific and Informational Approaches to Consciousness

Though the study of consciousness was traditionally thought to be within the remit of philosophy [26, 96], neuroscience is making progress in understanding the neurological parameters of conscious experience [21, 106, 107]. There have been multiple approaches to solving the ‘hard problem’ of consciousness; namely, how material processes can give rise to subjective, phenomenal experience [14, 88, 106].

One approach involves making observations of human brain activity while awake, asleep, comatose, under anaesthetic, and during psychedelic experiences to inform quantitative approximations of consciousness [106]. Other research has focused on certain modalities of subjective experience, such as visual perception and processing [21]. Other theorists have unbound consciousness from neural mechanisms, suggesting instead that it is a property of all systems with ‘integrated information’, including technological systems [125]. Many of these approaches have discussed technologies such as Alexa [39, 64, 94] and GPT-3 [16, 64, 106] in the context of consciousness, suggesting that these systems are important case studies in the debate. Yet, determining whether these systems, and more advanced future interactive technologies, can sustain conscious experience or not is still beyond the capacities of leading neuroscientific theories.

2.2 Perceptions of Consciousness

Bypassing the ‘hard problem’ of consciousness and assessing the ontological possibility of machine consciousness, other research has focused on understanding peoples’ perceptions of consciousness. The study of non-domain experts’ understandings of consciousness is also referred to as ‘folk theories’, ‘commonsense psychology’, and ‘heterophenomenology’ of consciousness [52]. Interestingly, much of the research into the perception of consciousness has used robots and cyborgs as exemplary edge cases in this debate, presented to participants to extract their understandings of consciousness in unfamiliar scenarios [38, 52, 119]. Gray et al. [38] explored two dimensions of the perception of mind: *agency* and *experience*. Using this model, they found that a social robot scored low on experience, and average on agency, when compared to humans at different ages and capacities, a selection of animals, and God [38]. Huebner [52] introduced participants to four hypothetical entities; a human with a human brain; a human with a CPU; a robot with a human brain; and a robot with a CPU. The results showed while participants were more likely to ascribe pain to humans rather than robots or cyborgs, there were no statistically significant differences in their ascription of beliefs to humans, cyborgs, or robots. This suggests that layman can be noncommittal or ambivalent in their perceptions of machine consciousness (PMC) within these edge-case entities.

Huebner [52] argues that these technological edge cases can expose the heuristics that people use to approximate consciousness in those other than themselves. However, this work was focused on informing the philosophical debate distinguishing between phenomenal and non-phenomenal states of consciousness, rather than understanding peoples’ conceptualisations of (technological) consciousness in their own right. Hence, the work only measured whether participants ascribed beliefs or the experience of pain to hypothetical cyborgs and robots, rather than explore their full, self-generated conceptualisations of consciousness across these entities. Also neglected in this work is how an individual’s PMC may influence their affective response to a technology and the subsequent interaction. Furthermore, there is little consideration of the consequences of PMC for the design of interactive systems. Finally, there has been significant developments in the capabilities of robots and artificial intelligence in the 12 years that have passed since the publication of Huebner’s results [52], and peoples’ perceptions of PMC may have evolved accordingly.

In the next section, we will describe the previous work in HCI exploring peoples’ perception of consciousness and adjacent concepts in interactive systems, and their impact on interaction. To motivate this, we first discuss the importance of understanding user perceptions and experience in HCI. We then describe the results of a literature review of 28 papers in the ACM Digital Library that mention ‘machine consciousness’. Finally, we briefly describe HCI research that has explored anthropomorphism and agency, concepts adjacent to consciousness, and how they may affect interaction.

2.3 The Importance of Understanding Subjective Experience in HCI

Shifting focus from the technical and human factors to experiential parameters of interaction has been an important development in HCI. While technical components clearly impact user experience (UX), users rarely understand how they work in fine or comprehensive detail. In absence of direct observation or clarity, users build mental models of how a technology works [20, 112], and form impressions and attitudes of its purpose and value [65]. These perceptions are shaped by environmental, cultural [67], contextual [7, 112] and further individual factors [10, 27]. Rather than understand the ontological nature of technological features themselves, empirical research into user experience attempts to collect and articulate subjective perspectives of technologies and their features.

This empirical approach has been important for understanding the different qualities that people perceive in technologies. This includes the pragmatic, such as safety [112], explainability [51], and usability [86], the hedonic, such as joy, delight and companionship [43, 87], and the eudaimonic, relating to technologies that facilitate self-actualisation and progress against personal goals [80]. In this work, we continue the focus on subjective experiences in HCI, by studying peoples’ perceptions of (machine) consciousness and how this perception may influence interaction with technologies such as GPT-3, Alexa, and vacuum robots. By doing so, we aim to explore HCI’s role in the current interdisciplinary problem of understanding machine consciousness. Further, in line with the tradition of imagining and creating technological futures [99], we see

the coming era of (potential) artificial consciousness as a challenge for HCI which requires addressing now.

2.4 Perceptions of Machine Consciousness in HCI

To review work on perceived machine consciousness in HCI, we searched The ACM Full-Text Collection (666,115 records), for the following terms [All: "machine consciousness"] OR [All: "technology consciousness"] OR [All: "artificial consciousness"] OR [All: "machine awareness"] OR [All: "self-aware technology"].

Our search resulted in 28 papers. We excluded one paper not written in English [76]. Three further papers were excluded, one because it had been retracted [123], and two others because the authors used the term 'consciousness' in an alternative way, either in a figurative sense [25] or to refer to peoples' awareness of technologies available to them [91]. We categorised the remaining 24 papers to identify key domains of previous machine consciousness research in HCI. We identified five categories; Developing Awareness in Machines ($n = 14$), Critical Approaches to Machine Consciousness ($n = 6$), Attitudes Towards Machines ($n = 3$), and Importing Theory ($n = 1$). The first category, Developing Awareness in Machines, includes papers concerned with technical approaches and problems in developing machines that can be context and content aware [19, 23, 48, 61, 62, 66, 132], self-aware [24, 28, 53], or demonstrate cognitive capacities such as thinking, learning [40, 131], explaining [1], or using language [90]. The second category, Critical Approaches to Machine Consciousness, includes papers that discuss either the abstract concept of machine consciousness [103], the broad risks and challenges of machine intelligence [63], or critique current approaches to understanding machine consciousness [50, 93, 114]. Attitudes Towards Machines is the category most relevant to this paper, and included papers on the demand for more context-aware and self-aware machines in a manufacturing context [137], broader fears of intelligent robots [128], and the perception of computer-generated teammates in gameplay [129]. The final category included a paper that imports neuroscientific and philosophical theories to support development of machine consciousness [71].

This review shows that there is literature theorising the technical means to achieve machine awareness, critical reflections on the concept, and the attitudes towards machine consciousness in three specific contexts. However, none of these papers directly address or theorise on peoples' perceptions of machine consciousness nor the consequences of this for interaction. In this paper we explore commonsense conceptualisations of (machine) consciousness in an online survey and contextualise our findings with relevant work from HCI and beyond, to address this gap.

2.4.1 Adjacent and Related Concepts to PMC in HCI. Consciousness is a heavily conflated term, so to be comprehensive we will describe HCI research that explores concepts adjacent to consciousness. As consciousness is perceived subjectively, it is intuitive to draw upon one's own consciousness to understand the consciousness of others. Because of this, we will briefly account the role of anthropomorphic features in ascribing parameters such as emotionality, intelligence, and self-awareness to machines. However,

non-human-like objects can also be ascribed consciousness, particularly through the attribution of an intentional stance [26]. To address this, we make an account of the literature concerning machine intentionality and agency in HCI.

Anthropomorphism and PMC. Peoples' subjective experience of consciousness is mediated by the fact that they are human. Therefore, the perception of technological consciousness may be likely to occur for systems and devices that are perceived as 'human-like' and are susceptible to anthropomorphism [33, 41, 94]. 'Human-like' features can consist of physical traits, such as facial features, body configuration, and gait, as well as cognitive capacities often perceived as human-specific, such as language, intelligence, and emotions [42, 138]. While 'Uniquely Human' traits refer to higher cognition, intelligence, and refinement, 'Human Nature' is instead associated with warmth, emotionality, desires and openness [34, 42]. All these traits could be associated with having mental states and consciousness. In HCI, these features have been particularly important in the design of robots [29] and chatbots [68]. Złotowski et al. [138] found that participants' perception of 'human likeness' in a robot was more affected by perceived emotionality than perceived intelligence, and can also mediate levels of empathy, user trust and relationship development with voice assistants [54, 98, 109].

A particularly distinct anthropomorphic feature utilised in many interactive systems is language. Voice assistants, such as Alexa and Siri, are able to detect human voices, extract meaning from them, and respond in spoken language with questions and answers. Natural language models (NLMs) have progressed significantly in recent years, evident in systems such as OpenAI's GPT-3 [31]. These language models can even make explicit claims of self-consciousness and awareness, which many, including developers of the software, find convincing [82, 120, 124]. However, these sophisticated models can also make emphatic claims to *not* being conscious, or to being a squirrel, a T-rex, a vacuum cleaner, or the Chicago River [111]. This suggests that a computer self-proclaiming to have self-awareness can be misleading, when they can proclaim to anything when asked.

Object Agency and PMC. Anthropomorphic traits may be misleading in terms of perceiving consciousness, which is not always understood to be a human-specific trait [8]. Even without human-like physical or cognitive traits, machines may be perceived as having subjectivity or consciousness, made apparent through the way they move and act. Humans have a common tendency to attribute agency and intentionality to objects [47, 84, 104]. While the term agency can refer to the 'capacity to exercise control over processes, motivation, action, and environment' [55], the term intentionality instead refers to the presence of mental states that direct and coordinate behaviour [26]. When shown a film depicting geometric objects moving in particular directions and distinctive trajectories, people usually explain what happened using mental states [47]. They will use phrases such as "the triangle *wants* to...", and even build elaborate narratives explaining the shapes' behaviour. There seems to be an automatic instinct to interpret moving things as having agency and causality [84, 104]. This perceptual tendency also appears in the interpretation of the behaviour of interactive systems, such as automatic vacuum cleaners [87, 121]. Intentionality has also been ascribed to non-robotic systems. Users of personal music recommendation algorithms, such as those available on Spotify,

have attributed a theory of mind to the recommendation algorithm, and even express hypothetical dialogues with them [36].

Work within Human Robot Interaction (HRI) has investigated the role of Theory of Mind (ToM) in the perception of robot action [46, 49, 60]. ToM describes the ability to infer the thoughts of others or to ‘mindread’ [46]. While much of this work has focused on building a model of ToM within machines [49], Hegel et al.’s [46] described how a robot’s shape and embodiment can influence users’ tendency to attribute theory of intention or will to their actions. Ono et al. [89] demonstrated that a ‘mindreading’ system can help users predict robotic behaviour and understand unclear utterances. While this research explores interesting effects of perceived ToM on HRI, it does not explore peoples’ own conceptualisations of or attitudes towards machine minds. Indeed, Theory of Mind is a specialist term; in our survey of 100 people, describe in the following section, not one participant referred to Theory of Mind.

This work on anthropomorphism and agency in interaction raises interesting points in regards to the perception of machine consciousness (PMC). However, there has not yet been a dedicated account of peoples’ PMC in their own right. The need for an understanding of users’ experiences with pseudo-sentience has been acknowledged before [117]. In this study, we take an explicit focus on peoples’ conceptualisations of machine consciousness. Firstly, we were interested in if PMC occurred, and how commonly. Secondly, we were interested in the concerns, curiosities and narratives they associated with machine consciousness. We identify dynamic tensions within peoples’ PMC, and suggest how these can inform challenges for HCI.

3 METHOD

To explore user perceptions of machine consciousness (PMC), we administered an online survey design in Qualtrics with 100 participants using Amazon Mechanical Turk (MTurk). We wanted our survey to establish a) whether PMC is a real phenomena, and b) an early picture of the narratives and concepts associated with PMC. First, we explored participants’ general conceptualisations of consciousness and machine consciousness, respectively. Then, we introduced participants to an example technology, and asked further questions related to this system.

3.1 Participants

We created the questionnaire using the Qualtrics survey platform and recruited participants via Amazon Mechanical Turk. As we were interested in qualifying the existence of perception of machine consciousness, surveying a large number of participants from diverse backgrounds seemed an appropriate approach. MTurk respondents remain anonymous, so they can be honest in admitting their perceptions, and free them from pressures to align with any expectations they perceive from researchers [44], but see [3]). We hope to use this online survey to assess whether PMC is a real phenomena, before future research embarks on substantial qualitative assessments of the phenomena via interviews or ethnography [65].

Due to incomplete or clearly automated responses, we rejected 33 responses. The final sample was 100 participants. The distribution between different example technologies was as follows: Alexa—30 respondents, GPT-3—29 respondents, and the vacuum robot—41

respondents. The regional distribution covered the US, the UK, Europe and Australia. Responses were gathered over a period of three days. Participants had to have an approval rate of at least 95% and over 1000 approved responses. They received \$1.20 in compensation for 10 minutes of work.

Participants reviewed an information briefing and gave their consent before beginning the study. The study was exempt from an ethics review as per local regulations of the main institution. Participants were given the chance to submit their email address if they wanted to be informed about the results of the study. Apart from this email address, no personal data was collected nor stored in this study. All data was stored on an encrypted device, only accessible by the authors.

3.2 Selecting and Portraying Example Technologies

Our first research question was whether people can perceive consciousness in currently available technological systems. The features of particular systems are likely to influence these perceptions, so it was important to include a variety of technologies as research probes. To identify and choose these example systems, the authors explored literature discussing consciousness generally, as well as machine consciousness in particular. We also considered the prevalence of these technologies in HCI literature. Subsequently, we chose three systems which have been the subject of debates on consciousness and of research in HCI: virtual assistant technologies, e.g. Amazon’s Alexa [39, 64, 94]; chatbots, e.g. OpenAI’s natural language model, GPT-3 [16, 64, 106]; and a vacuum robot [22, 87, 121]. These systems have regularly sparked speculation about whether they may be conscious or not, but there is no evidence whether most people entertain this or why. Hence, these systems seemed appropriate research probes for an investigation into PMC.

As discussed in the previous section, many features and behaviours have been associated with consciousness and adjacent concepts; our example technologies demonstrate varying levels of interactivity [55], anthropomorphism [42], language skills [103], and movement [47, 104]. For example, while the vacuum robot can move independently, the Alexa speaker and GPT-3 can use language and converse; while GPT-3 interacts via a text interface, Alexa uses a voice interface. This diversity allows us to probe whether PMC is observable across systems with a wide range of functionalities and traits. Future work could choose systems with different features, and isolate the complex and connected contributions of these features within PMC e.g. comparing PMC between GPT-3 with a voice interface and without.

While all three example systems are available to the public, it was crucial to explore technologies that participants are potentially familiar with (vacuum robot and Alexa) as well as those that are less well-known but still demonstrable within a short video clip (GPT-3)¹. This variety is important because the perception of consciousness may feasibly change with increasing awareness and interaction.

To introduce participants to these systems, we created or sourced three demonstration videos to introduce Alexa [105], the vacuum

¹This study was completed before the popular release of Chat GPT-3 in November 2022.

robot [126], and GPT-3 (video created by the authors and available in the auxiliary materials). Video vignettes and concept videos are regularly used in interaction analysis, speculative research and technology acceptability studies [56, 65, 135]. For participants, video demonstrations can be more engaging and immersive than written descriptions of technologies [74, 115]. They also offer higher control when compared to live demonstrations or self-directed interaction [65]. The videos were 34–71 seconds long; short descriptions of the contents of the videos can be found in Table 1. The videos contain unavoidably different portrayals of the different technologies, as they have varying capacities and features e.g. in contrast to Alexa or GPT-3, the Vacuum robot cannot use language but it can move independently. These videos were intended to be research probes for speculation, rather than independent variables; as a result, we do not make direct comparison between the example systems in our analysis. Our method presents a starting point for the assessment of PMC — future work may take a comparative approach and require greater standardisation between conditions.

3.3 Survey Design

The survey was designed using Qualtrics, and consisted of 15 questions across five sections. Participants took an average of eight minutes to complete the survey. The full survey is available in the auxiliary material.

First, participants entered their MTurk ID, read a brief neutral introduction to the survey and gave their informed consent. Then, in random order, participants were asked four questions about their general understanding of consciousness. There is little known about layman interpretations of the word ‘consciousness’ [119], let alone of machine consciousness [52, 117]. It was therefore important to establish how standard interpretations of consciousness can vary, and then see how these compare to conceptualisations of machine consciousness. Understanding peoples’ prior conceptualisations of consciousness could inform the design of future technologies that may be perceived as conscious.

One of these questions simply asked for participants to describe when they would use the word ‘conscious’. The other three were sentence completion tasks, which asked the participant to finish a sentence. Sentence completion tasks facilitate more flexible responses than questions with multiple choice answers, allowing respondents to use their own words and concepts in response to a prompt [69]. Sentence completion tasks can also guide the respondent towards an appropriate concept more clearly than open-ended questions, and are quicker to analyse for a large sample size.

Participants completed the following sentences: “I know something is conscious when it...”; “When I am conscious, I am...”; and “To me, the word ‘consciousness’ means...” This range of prompts was used to explore the various ways in which the words ‘conscious’ or ‘consciousness’ could be used. What is one’s own consciousness like? What could the word ‘conscious’ mean by itself, when not attached to an entity? There is some redundancy between the prompts; it has been demonstrated that similar questions can encourage different answers and richer insights [69]. This was important because we wanted to explore the wide range of meaning that is attributed to the term ‘conscious’, from sentence to applied

effort and focus, or other interpretations that we may not have been aware of.

The second section of the survey focused more precisely on interpretations of ‘technological consciousness’. In the survey, we used the term ‘technological consciousness’ rather than ‘machine consciousness’. While the term ‘machine’ is associated with mechanical apparatus, perhaps with industrial connotations, the term ‘technology’ applies more generally to ‘applied scientific knowledge for practical purposes’ [12]. Also, the word ‘technology’ appears to be in wider use in contemporary texts [127].

The first question explicitly asked participants to consider the necessary criteria for a technology to *appear* conscious. This is aligned closely to our interest in *perceived* machine consciousness. As the previous section focused on consciousness purely, here we wanted participants to consider the features that could make a technology *appear* conscious i.e. behaviours, interactions, abilities.

The second question asked participants to make a direct comparison between human and technological consciousness. Lallemand et al. [69] found that comparative sentence completion prompts can extract some of the most insightful responses. In our case, it is relevant to compare human consciousness, which our participants have first-person experience of, and technological consciousness, the subject of our inquiry. In their answers, we were interested to see whether our participants cited quantitative or qualitative differences between the two, as well as the valence of these responses: is technological consciousness considered better or worse? These comparative attitudes could influence PMC, and inform the interaction design of technologies that may be perceived as conscious.

In the next section of the survey, participants were randomly assigned to watch one of three videos introducing the interactive technologies, i.e. a robot vacuum cleaner, GPT-3 or Alexa.

As the example technologies are already available for use, some participants may have familiarity with them and pre-existing attitudes towards them. Existing relationships and attitudes can influence user experience of a technology [58, 65]. To assess whether these attitudes could impact our participants’ PMC for the example technology, we asked them to indicate their familiarity with it. To do so, they could choose one of the following options: ‘I have used this technology extensively’; ‘I am aware of this technology but have never used it’; or ‘This is the first time I have heard of this technology’. This question allows an initial assessment of the effect of familiarity on PMC; building a comprehensive account of how PMC may change with time and interaction is an important future direction for research.

As well as qualitative conceptualisations of machine consciousness, quantitative assessments may also be made. To our knowledge, there are no validated instruments to assess the perception of machine consciousness. While some theories of consciousness suggest it is all-or-nothing trait [85], some neuroscientists have suggested that there may be a spectrum of consciousness, spanning from death, through anaesthesia, dreaming, awake, to psychedelic experiences [106]. To explore the former, we asked participants to answer the following questions in regards to the example technology:

“Do you think it ‘feels like something’ to be this technology?”

Table 1: Descriptions of video demonstration for each example technology

Example Technology	Video Description	Attribution
Alexa	In a dark studio space, the small smart speaker initiates a conversation with a user about fashion and tattoos. The smart speaker asks whether the user has a tattoo and whether it hurt when they got it. It claims it was worried that it had hurt, claims to have a tattoo of the Amazon logo, and compliments the users' personality. The user and others listening in on the conversation are laughing and smiling at the interaction.	[105]
GPT-3	A screen recording shows a user typing a greeting and asking what the bot thinks about nuclear power. The GPT-3 chatbot writes that there are many opinions on nuclear power, but that it thinks that it has clean and efficient attributes. The user asks three more questions about building nuclear power plants, to which GPT-3 responds with relatively formal, one-sentence answers.	
Vacuum Robot	The white vacuum robot moves along the floor, with a blinking red light and a rotating sweep. The vacuum robot moves into a room where a person is sitting on a sofa on their phone. This person does not interact with the vacuum robot, which at first moves towards the sofa, then turns and moves away.	[126]

Participants could answer this on a five-point Likert scale, ranging from 'Definitely not', via 'Might or might not', to 'Definitely yes'. They could also select 'I do not understand the question'. The phrasing was inspired by Nagel's famous thought experiment [85], where describing 'what it is like' to be something can explore whether that thing has subjective conscious experience. This data can show whether this influential way of thinking about consciousness can be interpreted by non-experts, and applied to technologies. If so, this phrasing could be useful in future studies exploring PMC.

The following question asked participants to place the technology on a 'scale of consciousness'. They could indicate their answer on 11-point Likert scale, where zero equals 'Not conscious at all', and 10 equals 'Highly conscious'. Previous research has shown that zero is associated with nothingness or the perception of absence [5]. We presented the choice on a visual analogue slider to reduce bias [77]. This quantitative scale allows participants to give a binary response (totally unconscious or completely conscious), as well as a graded response (partly conscious). Using this question, we could calculate the average 'degree of consciousness' perceived in the technologies, if at all. The purpose of this question was not to establish a consistent spectrum of consciousness across participants, but to see whether participants could perceive *shades* of consciousness, rather than binary states.

After completing these quantitative assessments, participants were asked to consider a scenario in which their example technology *did* appear conscious. Asking participants to speculate about hypothetical scenarios and design fiction is a well established approach in HCI, used for for evaluating 'upstream technologies', alternative futures, and for opening up new areas of research [32, 74]. The first questions asked them to describe their personal response if the technology were to 'appear conscious' i.e. Would it matter to them? How would it make them feel? Emotions and moods can have important influences on interaction, so understanding the affective components and consequences of PMC is a critical research question [79, 112, 136]. This initial sample of affective associations with PMC can inform the design of sensitive and appropriate devices that may appear conscious.

In addition to their emotional perceptions, we wanted to understand our participants' understanding of the hypothetical risks and opportunities of machines appearing conscious. Inviting consideration of both positive and negative implications can encourage a more nuanced discussion of technology experiences [74, 97]. We kept the questions broad and exploratory, so as to identify as many different perspectives as possible. Future work could build on this by developing more targeted methodologies for particular groups of people e.g. those who perceive machine consciousness versus those who don't, or people who are enthusiastic about machine consciousness versus those who are more hesitant or afraid.

3.4 Limitations and Scope

This study uses an online survey to sample peoples' perceptions and attitudes towards machine consciousness. The use of sentence completion prompts encourages participants to respond in ways that are meaningful to them, while also reducing the time required for completion and analysis. However, participants often give one- or two-word phrases in response to the prompts. This lightweight method is unlikely to capture the full richness of peoples' conceptualisations of machine consciousness, which may change over time and context. Our understanding of PMC would be further enriched by ethnographic methods attending to implicit, informal and embodied interactions. While there may be a risk of over-interpreting our results, they do provide significant evidence for the existence of PMC. This offers the robust motivation necessary for future work, which could use other methods to explore PMC across different contexts, and over extended periods of time.

Amazon's Mechanical Turk is widely used in academic research, and has been shown to provide reliable results in different contexts [44, 57]. However, there are some common concerns with the use of this software for recruitment. MTurk workers often complete many studies a day, while trying to optimise their hours-worked to payment ratio [44]. This may lead participants to give low-effort or highly efficient responses, rather than those which best represent their perceptions of consciousness. While we did receive many well-considered and meaningful answers, future work could use more

intensive qualitative methods, such as interviews and ethnographic inquiry.

In addition, our participants were mainly from the United States and Europe, and so may not be representative of PMC worldwide. For example, in certain cultures there is already a tradition of attributing presence to inanimate objects, such as *'kami'* in the Shinto belief system [108]. We also required a high level of English proficiency for this experiment; the meaning of 'consciousness' may change with translation. While this paper presents an initial survey of PMC, future work could take this inquiry into other cultures and languages.

MTurk workers necessarily have high computer literacy, with access to the internet and payment systems. Hence, their perceptions of technologies, and machine consciousness specifically, may be skewed due to this. In addition, MTurk participants may have a peculiar relationship with advanced computational systems. We rejected 33 responses to our survey, largely due to clear indications of automation and the use of phrasing taken directly from the internet. This suggests that many MTurk users are aware of the potential to deploy machines to emulate their task performance, and as a result, may have different conceptualisations of the opportunities for machines that can appear conscious (such as GPT-3). These limitations further highlight the need to build on the results presented in this paper with diverse participant samples.

The example technologies of GPT-3, Alexa and the vacuum robot were chosen due to their association with debates on machine consciousness and related concepts. However, there are many other technologies that may appear conscious that are either available currently or in development, such as android robots [73], magnetic slime robots [118], and artificial art generators, such as DALL-E 2 [17, 75, 95]. The rapid development of new technologies and their presence in the media could quickly influence peoples' perceptions of machine consciousness. We suggest that this work is repeated in the future to track and monitor changing attitudes towards machines that may appear conscious. Nevertheless, our inquiry provides an initial survey on peoples' relationships with PMC as it currently stands.

4 RESULTS

Here, we first present a quantitative analysis of the results of our survey. Next, we analyse qualitative responses to the survey prompts. We conducted quantitative analyses using non-parametric methods as all scores were either binary or ordinal data. Qualitative analysis was conducted using affinity diagramming [72]. Three researchers structured the responses to the open-ended survey questions into clusters which represented recurring patterns in participant answers. Given the large diversity in the data and moderate volume, we decided to adopt this pragmatic approach as suggested by Blandford et al. [9]. We account recurring themes and concepts with direct quotes and paraphrases from our participant data. In the case that more than five participants mentioned a particular theme, we have provided the exact number of participants instead of their individual participant IDs. To see each participant's numerical rating for their example technology, refer to Table 2, 3, and 4. The full dataset is available in the auxiliary materials.

4.1 Ascribing Consciousness to Technologies

We first investigate how many participants ascribed any consciousness to their example technology. As participants were asked to rank the consciousness of the technologies on an 11-point Likert scale from 0 to 10, we interpreted the responses '0' and '1' as a lack of perceived consciousness. We decided to include the '1' response as some participants may have interpreted '1' as an indication of having answered the question. This is partly due to the limitations of the survey tool—the question was presented as an analogue slider and moving the slider was required to ensure that the participants answered the question.

Figure 2 shows that more than half of participants perceived the example systems as showing a certain degree of consciousness. While 22 out of 41 participants perceived the vacuum robot as having some level of consciousness, 21 out of 29 participants did the same for GPT-3, and 17 out of 30 for Alexa. Overall, 63 out of 100 participants attributed consciousness to the example system they were shown. We conducted a one-way ART-ANOVA [134] to determine if the system presented had an effect on whether the system was perceived as conscious. We found no significant effect, $F(2, 97) = 0.96, p = .39$. Additionally, we checked if prior familiarity with the system had an effect on ascribing consciousness to it. An ART-ANOVA showed no significant effect, $F(2, 97) = 0.63, p = .53$.

In response to the question asking whether it 'feels like something' to be that technology, only three participants indicated that they did not understand the question. The most common answers were found at the extremes; 30 participants selected 'Definitely not' for their technology, while 27 selected 'Definitely yes'. This suggests that the majority of our participants could understand the phrasing used in Nagel's thought experiment [85] and could apply it to technological systems.

4.2 Degree of Consciousness in Different Technologies

Next, we investigate the reported degrees of consciousness among the participants who did not state that the systems fully lacked consciousness. Figure 3 shows how the scores were distributed in the full sample and in the three systems respectively. Using a one-way ART-ANOVA we checked for the effect of the system presented on the perceived degree of machine consciousness and found no significant effect, $F(2, 60) = 0.47, p = .63$.

4.3 Perceptions of Consciousness

We first analysed the participants' perceptions of consciousness. Participants associated being 'aware' (56 participants) of some *internal or external context*, including awareness of *the self* (e.g. P13, P19, 23) or *'memories, feelings, sensations, and environments'* (P44). Many participants emphasised the capacity for *'feelings'* (10 participants), *'sensation'* (P44, P88) and *'emotions'* (14 participants) as core elements of consciousness. Something that is conscious may *'feel everything like pain and also joy'* (P98), or may be simply *'sentient'* (P64, P71, P73, P86, P100).

Another often mentioned aspect of consciousness was morals (P34, P35). A conscious entity was often seen as one that follows good principles (P32, P33). Finally, some participants attributed a clear moral agenda to a conscious being as shown in this example




Figure 2: The share of participants who perceive a degree of consciousness in the three interactive technologies examined. We interpreted the ‘0’ and ‘1’ as a lack of perceived consciousness.




Figure 3: The distribution of perceived degrees of consciousness for all three systems examined in the study (top) and for each of the systems (bottom). Ratings of 0 or 1 are excluded.

where a conscious being was to ‘*protect your own interests and those of others*’ (P97), or to ‘*make decisions based on right and wrong*’ (P15). Furthermore, taking ‘*responsibility*’ for the actions one takes was a recurring theme associated with consciousness (6 participants). Consciousness was also associated with the presence of ‘*deeper meaning and a clear purpose*’ (P13).

Others associated ‘*reacting*’ to the environment (12 participants) and ‘*responding to stimuli*’ (6 participants) with being conscious. Some thought these responses ought to be rapid (‘*reacts fast*’) (P62),

and appropriate—‘*intelligently responds to stimuli*’ (P4). Two participants stated that something is conscious if it can ‘*consider concepts higher than impulse*’ (P72) or ‘*instinctive reaction*’ (P58), and it can demonstrate ‘*logic and awareness*’ (P58). People also associated consciousness with an ability to ‘*think for oneself*’ (P15, P22, P28, P37) and evidence of ‘*act[ing] under [one’s] own volition*’ (P8).

4.4 Perceptions of Machine Consciousness

We then asked participants to imagine the qualities of a conscious machine. Many participants were sceptical of such a concept insisting that machine consciousness would be *'impossible to truly attain'* (P4, P55, P76), *'unbelievable'* (P42), or being the product of imitation or deception (P5, P58, P72). Others suggested machine consciousness be *'artificial'* (P17, P35, P47, P58, P91), *'fake'* (P35), or *'robotic and too logical'* (P46).

However, other participants speculated about machine consciousness, and the criteria it must meet to be perceived. The first property which was often mentioned was describing the entity as *'being alive'* (P4, P24, P95). This was often associated with the technology's ability to *'to take action'* (P90, P47), *'react with its environment'* (P93, P8), be *'reactive to stimuli'* (P8), and *'in a way that living beings do'* (P2).

Others emphasised the needs for a conscious machine to *'understand and feel what is happening around it [...]'* (P21). Others suggested that a conscious machine would be *'aware of itself'* (6 participants) and of *'its own existence and condition'* (P95).

Further, some participants noted that a conscious technology would need to experience *'genuine feelings'* (P28), both of its own (P3, P29, P83, P62, P100) and understanding those of others (P21). A conscious technology would need to be *'able to feel love and hurt'* (P28). Feelings were also often mentioned together with morals and conscious machines were required to *'have moral understanding'* (P33, P34, P98) and *'distinguish right from wrong'* (P15). Another cited that *'making small mistakes'* (P98) and *'being stuck in moral dilemmas'* (P98) were necessary criteria for a technology to be perceived as conscious.

Another property of conscious machines was independence of mind, the ability of the technology to *'think for itself'* or *'on its own'* (7 participants), to *'take actions on their own accord without influence of their makers'* (P47, P90), and *'make independent decisions'* (P41, P64, P83). A conscious technology should *'understand human nature beyond its programming'* (P6) or *'instructions'* (P76), to be *'flexible'* (P14) and *'constantly learn and adjust its own programming'* (P37). Related to this, was the need for a machine to appear as *'alive'* (P4, P95) and *'to not act in a mechanical, automatic way'* (P95).

Other participants stated that conscious technologies would be *'similar to humans'* (P17, P68) or have *'human-like'* (P30, P35, P48, P75) properties as essential for a technology to appear conscious. This could include being *'able to feel emotions like humans do'* (P3), *'successfully act as a human'* (P56) or *'mimic human behaviour to the point where you are not aware it's NOT human'* (P86, P71, P72).

4.5 Machine vs Human Consciousness

We asked participants to directly contrast human and machine consciousness. Some participants referred to machine consciousness as *'primitive'* (P3) or *'limited'* (P21, P29), or being lesser, either in general (P50, P63, P82) or in regards to a specific concept: *thinking or understanding* (P84, P85, P93), *power* (P64), or *importance* (P28, P33). In contrast, some participants referred to machine consciousness as *'more advanced'* (P27), *'more useful'* (P34), or *'more powerful'* (P77). Others said that machine consciousness would be very different, or even *'incomparable'* (P32) to that of humans. One

participant suggested that machine consciousness could be *'to[o] different for us to notice it'* (P68).

Another participant remarked on how decisions taken by a machine would always be deterministic: *'Different as the way it makes decisions is informed and always following a certain roadmap or decision tree'* (P83), or *'governed by a computer chip'* (P88), and not bound by biological principles: *'a logical creation that does not abide by the infinitely more delicate biological processes'* (P65). Alternatively, other participants emphasised that a conscious machine would have to be *'aware that the rules by which it is governed are being followed or broken'* (P18).

A group of participants used this answer field to share their doubts about the possibility of achieving machine consciousness. Some remarked that machine consciousness was *'not possible as it doesn't have feelings'* (P55), yet others remarked that achieving technology consciousness would be *'an evolutionary leap'* (P37) and open new possibilities for consciousness such as eliminating the need for a bodily form: *'artificially created consciousness unbound by physiological boundaries'* (P54).

For some, machine consciousness was still regarded primarily as a man-made artefact which must be *'controlled by humans, so [it would be] artificial instead of real consciousness'* (P47). Artificial consciousness could also be a very elaborate tool, created and controlled by humans: *'unable to work without interference from a human being who would start and impact their consciousness (they would not be able to develop independently)'* (P90).

4.6 Reactions to Different Technologies

Using demonstration videos, we introduced the participants to one of the three following technologies; a vacuum robot (Table 2), GPT-3 (Table 3), or Alexa (Table 4). We asked them to speculate on the following, *if their example technology appeared conscious*: whether it would matter, whether they would care, and any risks or opportunities they perceived. Importantly, these answers may not relate to a participant's current perceptions of consciousness in the technology, but instead their hypothetical response if the technology did appear conscious. While there was a broad diversity of answers to these questions, the implications can be broadly categorised into the following themes:

- Performance implications
- Social implications
- Existential implications
- Absence of implications.

4.6.1 Performance Implications. Participants anticipated that the appearance of consciousness in machines could have implications for the machine's performance. Participants perceived that a conscious GPT-3 would become more able *'to do jobs that other people don't want to do'* (P48), and *'replicate humans on face-to-face tasks'* (P56). Both Alexa and GPT-3 may also come to more *'reasonable'* (P66, P85, P95), *'logical'* (P60) answers, be more *'helpful and informative'* (P55) as they *'actually know what [they're] talking about'* (P96). Participants anticipated that if the vacuum robot appeared conscious it might be *'useful to many'* (P27), by offering *'more optimal cleaning'* (P31) and *'clean rooms without the use of boundary tape'* (P20).

Table 2:
Vacuum Robot

ID	PMC Score
P01	00
P02	00
P03	00
P04	00
P05	08
P06	01
P07	09
P08	00
P09	00
P10	07
P11	10
P12	07
P13	00
P14	07
P15	04
P16	06
P17	00
P18	01
P19	05
P20	05
P21	06
P22	00
P23	08
P24	02
P25	08
P26	09
P27	04
P28	00
P29	00
P30	01
P31	01
P32	06
P33	07
P34	08
P35	02
P36	06
P37	00
P38	03
P39	06
P40	00
P41	01

Table 3: GPT-3

ID	PMC Score
P42	00
P43	08
P44	08
P45	09
P46	00
P47	00
P48	00
P49	07
P50	04
P51	06
P52	05
P53	00
P54	04
P55	04
P56	03
P57	07
P58	00
P59	08
P60	07
P61	05
P62	09
P63	06
P64	01
P65	02
P66	03
P67	02
P68	01
P69	06
P70	05

Table 4: Alexa

ID	PMC Score
P71	06
P72	00
P73	01
P74	07
P75	07
P76	00
P77	10
P78	08
P79	06
P80	08
P81	06
P82	05
P83	00
P84	04
P85	00
P86	00
P87	00
P88	01
P89	07
P90	02
P91	07
P92	01
P93	03
P94	10
P95	00
P96	04
P97	07
P98	07
P99	04
P100	00

However, participants also predicted negative implications for performance. They wrote that a conscious GPT-3 may give ‘*subjective answers*’ (P62), or become ‘*less reliable if acting on feelings*’ (P42). They imagined that a conscious Alexa could give an ‘*answer that makes you feel even worse*’ (P98). Participants worried that a conscious vacuum robot could ‘*behave badly*’ (P32), by causing intentional damage or harm, and ‘*suck different type of things that you don’t want [it] to*’ (P35).

4.6.2 Social Implications. As well as performance implications, participants suggested that PMC could have social consequences. The positive implications for a conscious Alexa including building ‘*better, fuller relationships with humans*’ (P95, P30) or becoming ‘*great at being a carer for older people living on their own, or people that have problems interacting with other people*’ (P86). Alexa may also come more context aware, by ‘*not interfering with people who don’t like it*’ (P97). One participant said they would ‘*love to have something to talk to that wasn’t family or friends*’ (P86, P9). Some anticipated that they would feel ‘*obliged*’ (P92) to a conscious Alexa, as it would ‘*deserve the respect accorded to conscious animals*’ (P92).

However, participants also perceived social risks, both for people and for the machines that appeared conscious. One participant was concerned that a conscious vacuum robot would be ‘*trying to talk to [them]*’ (P16), or alternatively would ‘*judge [them] or communicate private things about [them] to others*’ (P3). They also cited risks of ‘*deception*’ (P42), feeling ‘*spied upon*’ (P87), or ‘*that someone may think they are connecting with a human when it is not a human*’ (P55). Once also described the risk that GPT-3 could be ‘*manipulated to do wrong things*’ (P68). Perceiving affective implications, some participants said PMC could be ‘*creepy*’ (P56, P71, P95) and some may be ‘*scared*’ (P54, P68, P71, P73, P85) if Alexa appeared conscious. Some thought it was probable that humans would ‘*freak out... probably become destructive*’ (P100).

Others perceived moral risks for a conscious Alexa, who may not ‘*be afforded the same rights as other sentient beings*’ (P86). Participants asked whether it be ‘*okay for us to use AI the way we have if they are conscious*’ (P83), or saying that using a conscious Alexa ‘*would be similar to slavery*’ (P96). Another participants said that a conscious Alexa would face the ‘*same risks any individual faces: to grow up in an unhealthy environment, pick up negative traits and develop irrational sentiments, like hate, prejudices etc*’ (P95).

4.6.3 Existential Implications. Participants discussed the implications for mankind in an era of conscious machines, as well as for the existence of conscious machines themselves. One participant said that Alexa could ‘*even become some sort of benevolent entity that, thanks to her superior knowledge, can “guide” humanity*’ (P95), while GPT-3 could ‘*imagine things that might otherwise be impossible*’ (P61). For the technologies themselves, participants imagine they could experience self actualisation, being ‘*allowed to flourish on their own*’ (P83), and ‘*further the understanding of machine intelligence and promulgate the idea among ordinary people*’ (P92). One participant notes the opportunity for a conscious vacuum robot to represent ‘*human-originated life...after humans are all dead*’ (P8).

In contrast, some participants noted significant existential risks and fears of machines appearing conscious. A conscious vacuum robot was anticipated to ‘*revolt against its own creator*’ (P30), maybe becoming ‘*evil and harm humans*’ (P9). Participants wrote that a conscious Alexa could ‘*interfer[e] with our lives*’ (P84) or ‘*form its own intentions which are inimical to our plans*’ (P92). Several participants imagined the potential for a ‘*Terminator scenario*’ (P66) with conscious GPT-3, where ‘*robotic technology could overpower human thought and existence*’ (P48) and ‘*decide humans need to all die*’ (P46).

Some participants also empathised with the hypothetical conscious technologies, suggesting that GPT-3 ‘*might realise what it is*

and have some kind of crisis. This might make it seek out tasks that it is not designed to do and find a purpose for itself' (P61).

4.6.4 Absence of Implications. It is worth mentioning that a significant number of participants did not describe any risks (33 participants out of 100) or opportunities (32 participants out of 100) associated with machines appearing conscious. Sometimes this indifference was accompanied by a dismissive sceptical phrase, such as *'it is just a robot'* (P16, P18, P35, P84), or *'just an appliance'* (P38). Others say there were 'no risks because it is not possible' (P47, P48, P58, P89), or 'the concept [of consciousness] does not apply' (P13).

5 DYNAMIC TENSIONS IN THE PERCEPTION OF MACHINE CONSCIOUSNESS (PMC)

In our analysis of the data, we identified key themes evident within our participants' perceptions of machine consciousness (PMC). As consciousness is a heavily conflated and discursive term, we have found themes that overlap, diverge, and even contradict each other, as seen in Figure 4. In this section, we identify and discuss a series of dynamic tensions within PMC that arise from conceptual, pragmatic, and philosophical challenges and inconsistencies. These tensions are not mutually exclusive, and are likely to interact with and influence each other. We hope that they provoke further discussion, questioning and research into the phenomena of PMC from the HCI community, and present as a step towards a comprehensive theory of PMC. In the following section, we critically reflect upon the proposed tensions and explore the implications for the design of interactive systems that may be perceived as conscious.

5.1 Denial versus Speculation

Our quantitative analyses suggest that over half of the participants ascribed consciousness to the example systems of GPT-3, Alexa and a vacuum robot. However, the rest did not perceive any level of consciousness in these systems. In their responses, some participants thought that machine consciousness would be impossible to achieve for robots and technologies, and any associated traits could appear fake or artificial. This rejection of the principle of machine consciousness may prevent people from perceiving machine consciousness in systems or admitting to doing so. We could refer to this subset of participants as machine consciousness *sceptics*.

On the other hand, the majority of our participants were willing to ascribe consciousness to the example systems, and to entertain the idea of machine consciousness more broadly. This subset of participants could be referred to as machine consciousness *speculators*. Among these participants, there were very different conceptualisations of what machine consciousness could look like, and their subsequent attitudes. Many thought machine consciousness would have to be similar to, or the same as, human consciousness, while others thought it would be more technical, objective, or logical. Others thought machine consciousness would be incomparable to that of humans, while others predicted that it could be better or worse than human consciousness. It is these speculative conceptualisations that we explore further in the remaining dynamic tensions. However, the first tension here can be summarised as the contrasting denial and acceptance of the premise of machine consciousness by our participants. Future work may find further demographic categories beyond the sceptics and speculators identified here.

5.2 Thinking versus Feeling

Many participants proposed that thinking is critical to consciousness. These people perceive evidence of intelligence, logic, and considered decision-making as indicative of consciousness, or incoherence and lack of understanding as impediments to PMC. Thinking can be demonstrated by bypassing impulse and reflexes, in favour of careful consideration. In contrast, other participants placed emotions and feelings as central to consciousness. These people may perceive consciousness in machines who demonstrate feelings of pain, express emotions such as love or frustration, or are able to understand and interpret the emotions of others.

While these traits are not mutually exclusive, one took precedence over the other for many participants. Good decision making can require emotional intelligence or empathetic sensitivity in certain contexts. For example, a robot that corrects a child's spelling in a letter to their parents may have high literacy intelligence, but poor emotional perception. Alternatively, because emotions are not always logical or consistent, demonstrating feelings could undermine a machine's reputation as a rational and reasoning entity. Depending on the importance a user places on thinking or feeling as a criteria for consciousness, a spell-checking robot may be perceived as more or less conscious. Hence, our second dynamic tension in PMC is that found between the relative importance of thinking and feeling in PMC.

5.3 Interaction versus Experience

Some participants considered the interactivity of a technological entity to be a key indicator of its consciousness. Fast and appropriate reactions to stimuli and the capacity for action and movement could be taken as signs of machine consciousness by these participants. This aligns well with the 'AGENCY' model in folk psychology, where motion trajectories and contingent interaction can trigger automatic attribution of consciousness [4].

On the other hand, some participants described genuine awareness of self and surroundings as critical to PMC. Here, awareness was not tied to any corresponding expectations of machine functionality. Awareness is closely related to the concept of sentience, which is the feeling of sensation, distinct from perception or thought [81]. Awareness and sentience describe the existence of qualitative phenomenological experience, where it 'feels like something' to be that entity [85].

At first, the traits of experience and interactivity may at first seem to be inextricably linked. How can something demonstrate experience without some level of responsiveness, and how can something act without first experiencing a cue or input to respond to? However, these two traits can be tentatively separated. Current technologies have capacities for conditional interactivity; a motion sensor can respond to movement, but may not be described as having 'experience'. As the subjective experience of others is difficult, if not impossible, to prove, there could be a risk of putting all interactive behaviour down to rules and associations, and denying subjective experience unfairly. The risk of false-positive attribution of complex cognition has been heavily discussed by researchers in complex animal behaviour [113].



Figure 4: The five dynamic tensions are shown in their opposing categories. Key underlying concepts from the participant data characterise each category e.g. Denial summarises concepts such as “impossible”, “artificial and fake”, “not applicable”, “unbelievable”, and “just a robot”.

On the other side, some comatose patients and patients with ‘locked-in’ syndrome have subjective experience despite being unable to interact with the outside world [106]. These patients’ brain activity can be comparable to that of non-patients [107], yet they cannot physically respond to external stimuli. An analogous technological example could be a ‘conscious’ computer or GPU, that does not have the ability to display output on a screen; it may have experience but cannot express it. Hence, while the traits of experience and interactivity initially appear to be intrinsically linked, a more complex dynamic relationship hides between the two.

5.4 Controlled versus Independent

Some participants perceived machine consciousness as being ultimately controlled by and dependent on people. This could either be understood as machines being inevitably programmed by humans, or always submissive to human wishes and desires. This stands in contrast to other conceptualisations of consciousness, which were often associated with the ability to think and make decisions for oneself. Participants also mentioned that a conscious actor must also take responsibility for their actions. Before attributing consciousness to a machine, the participants may need evidence that a machine can think independently and have sovereignty over its own thoughts, actions, purpose and opinions.

As technologies are currently designed, built, and deployed by humans for their own will, for some people they may never achieve true self-determination. For people that prize independence from the authority and control of others, this could prevent them from perceiving machines as conscious. In this way, the tension between control and independence could motivate the denial of machine consciousness. However, other participants do not see independence as critical to consciousness, and even cite the risks of technological slavery, where a conscious machine could be subjected to work against their will. For these participants, denying machines the rights that other sentient beings are afforded could be a moral injunction. A further implication of the tension between control and independence is the status of machine learning (ML) as a form of ‘programming’. While a developer determines the input for black-box machine learning, the output is the product of abstract statistical associations, largely impenetrable to human understanding [102]. Even some experts see the capabilities of these computations as inherently magical or enchanted [13]. These interpretations of ML-derived intelligence could cause individuals to perceive AI as acting independently from human determination. In this case, an AI may be developed by humans, and put to use by humans, but their analyses and insights could be considered products of conscious, independent thought.

5.5 Rigidity versus Spontaneity

For some participants, consciousness was associated with rule-abiding diligence and carefulness. Participants placed prominence on individuals’ moral principles, in particular that of not harming others, as well as having a careful, ‘conscious’ approach to tasks. These suggest that maintaining rigid discipline, consistency, and moral integrity may be associated with the perception of consciousness in humans. Acknowledging the technological nature of machine consciousness, some participants perceived machine consciousness as being pre-defined and programmed. To these participants, machine consciousness would be carefully designed to follow rules and instructions, and showing rigid responses to stimuli.

However, other participants highlighted spontaneous, expressive and flexible action as an indicator of consciousness, in both people and machines. In machines, evidence of imperfection, making small mistakes, and breaking of rules could be considered criteria for PMC. If a machine could demonstrate an awareness of its rigid programming, and managing to transcend or break these rules, this could be a strong indication of consciousness in machines for some people.

While the previous dynamic tension described the relationship of consciousness with authority, control and independence, this tension instead demarcates the tension between the conscious effort required to stick to one’s rules and principles rigidly, and the creative spontaneity needed to break rules and act unpredictably. If principled behaviour is a conscious trait in humans, but spontaneous rule-breaking behaviour is associated with consciousness in machines, perhaps rigidity and spontaneity implicate these entities differently. Either way, it could be argued that to show unpredictable and rule-breaking behaviour, a baseline of normal rigid behaviour

and rules to break is first required. Hence, these two concepts are in dynamic tension.

6 DISCUSSION

While the prospect of machine consciousness has received increasing attention in recent years, there has been no structured, empirical account of peoples’ perceptions of machine consciousness (PMC). To address this gap, we asked 100 layman to rate three currently available interactive systems (GPT-3, Alexa and a vacuum robot) by level of consciousness. We also asked about their conceptualisations of consciousness, both organic and technological, and the risks and opportunities presented by machines that appear conscious. Through this, we demonstrated that a majority of participants perceive at least some level of consciousness in the example systems. Using the qualitative findings, we have extracted and formulated a series of dynamic tensions within peoples’ perceptions of consciousness. These dynamic tensions can be described in five pairs of opposing concepts; Denial versus Speculation, Thinking versus Feeling, Interaction versus Experience, Controlled versus Independent, and Rigidity versus Spontaneity. In this discussion section, we will discuss the opportunities and challenges afforded by our insights, both for theory and for the design of interactive systems in an era of PMC.

6.1 Initiating Theory of PMC through Dynamic Tensions

The five dynamic tensions presented in this work represent the thoughts and ideas collected from 100 participants in an online survey. They are presented as an initial description of the conceptual, pragmatic and philosophical inconsistencies and relationships that enrich PMC, acting as a springboard for further discussion, analysis and theory. We hope that future work will extend understanding of these tensions and identify new ones.

The dynamic tensions presented here are not thought to be mutually exclusive. They may mediate and influence each other. For example, if someone denies the existence of conscious machines (subsection 5.1), and prioritises emotionality as a core element of consciousness (subsection 5.2), a machine designed to show emotions could be perceived as manipulative, fraudulent, or even just bizarre. For a machine under the control of a malicious actor (subsection 5.4), their inability to break free from these chains be seen as an example of inflexible behaviour (subsection 5.5), and add further fuel to the denial of their consciousness (subsection 5.1).

These reactions and counter-reactions between design and PMC should be of critical interest to HCI scholar and interaction designers. Anticipating peoples’ attitudes to machines that appear conscious could influence the design and deployment of interactive systems. Next, we describe some challenges for HCI and interdisciplinary collaborations that may emerge from the dynamic tensions present in peoples’ perceptions of machine consciousness.

6.2 Challenges for HCI

While the dynamic tensions above describe complexities within the phenomena of PMC, the following challenges speculate about the resulting interactional and societal implications that may follow. The following list of challenges for HCI in an era of perceived

machine consciousness is by no means complete. Instead, they challenges are presented to highlight the wide-ranging consequences of PMC. While the challenges are certainly within the remit of HCI, the field will not be able to address them alone. These challenges require interdisciplinary efforts to build multilateral solutions across law, politics, industry, philosophy, and many other fields. However, while these fields may be more concerned with the ontological reality of machine consciousness, we instead speculate on PMC. With this stance, we dismiss the thorny ontological dilemmas of other fields, while speculating on the interactive challenges in technological futures where we share our world with machines that appear conscious [6]. The data used in this paper is available to the research community to facilitate further analysis and kick-start interdisciplinary discussion that can prepare us for such a world.

6.2.1 Hierarchy and (In)equality. Our results indicate that both sceptics and speculators of machine consciousness will be significant in number. With such distinct positions, we may expect a polarised and controversial discourse to emerge on the status of machines that are sometimes perceived as conscious. The implications for hierarchy and (in)equality between conscious entities present rich design spaces for HCI, who could build for a variety of positions on this issue. Some sceptics may show their conviction in their denial of machine consciousness by arguing against equal rights for such machines, or even showing violence towards them (*"I would never hurt something that could actually experience pain"*) [100]. Others may justify the use of machines because their consciousness is 'lesser' than that of humans, while others may lobby for equal rights for machines. This debate may also draw further attention to the status of other entities, such as animals and plants, in our society; perhaps an unexpected alliance between AI and livestock could emerge. There may be people who perceive machine consciousness as greater than that of humans, leading to respect, reverence, and perhaps (religious?) submission towards such machines. Politicians and leaders could leverage this to increase their authority. Rather than speculating whether machine consciousness is lesser than or better than human consciousness, HCI should scrutinise the implications of these respective positions for interaction and their relative ethical integrity.

6.2.2 Empathy and Relationship. Feelings and emotions were central to many of our participants' conceptualisations of consciousness. As a result, we may expect people to build empathetic attitudes and sincere relationships with machines that they perceive as conscious. Humans have capacity for an immeasurable number of types of relationship, including friendship, rivalry, romantic entanglements, and professional. We can expect similarly complex formulations of relationships between humans and the machines they perceive as conscious. Some people may refuse to use machines they perceive as conscious because they feel guilty or immoral for using or enslaving them. Other people may feel it necessary to reward machines for their work and effort, raising questions of how we build a system of distributing value that works for both humans and machines. People may perceive robo-therapists as needing occupational therapy and support [133]. If a conscious-appearing machine breaks or is destroyed, people may grieve for it, prosecute for machine-slaughter, or even seek revenge. The availability of machines that appear as conscious could also cause changes in

the perception of and demand for human-to-human relationships, which have different modalities. Depending on the sophistication of the machines, there may be a mismatch between users' expectations of their machines and their actual ability to provide emotionally intelligent support. Rather than focusing on how machines feel or the mechanisms of machine emotion, HCI should mitigate the consequences of empathetic liaisons between humans and machines, so that people are not let down emotionally by their machine companions (or vice versa).

6.2.3 Authorship and responsibility. Our findings show that consciousness is associated with an independence of thought and self-derived purpose. If machines are perceived as conscious, it may be that they are attributed authorship and responsibility for their outputs. When this output is deemed useful or beneficial to society, machines may have to be included in trademark regulations, licensing laws, and citation etiquette. Alternatively, when these outputs are deemed inappropriate, dangerous, or criminal, we may have to find ways to hold machines, and/or the people who use them, accountable for their actions. This challenge is not to create machines that can judge criminal behaviour, but to determine whether machines can be *judged to have criminal behaviour*. Without this, the ability of the legal system to address malicious actors may be questioned. In the past, the prosecution of non-human actors, such as dogs, asserted the importance of individual responsibility [122]; with the emergence of machines that can commit criminal acts, we may have to revisit our formulations of individual justice and retribution. Rather than wrestle with the origin of innovation and who or what should take the whole blame, HCI could focus on designing for the division of glory and the sharing of responsibility.

6.2.4 Stewarding the Perception of Machine Consciousness. Researchers hoping to develop 'real' machine consciousness should consider how this subjectivity will be portrayed in an interface. With a palette of traits that can lead to the perception of machine consciousness, interaction designers may be able to manipulate to what degree people perceive a system as conscious. There could be a scenario where designers are seen as denying a machine the chance to present as conscious, for example, through a lack of interactive, emotional, or spontaneous features. As explored above, there are major ethical, political and social consequences to the perception of machine consciousness, within which designers are directly implicated. HCI scholars and interaction designers are stewards of the perception of machine consciousness in an interface. Our responsibility is not to define the existence of machine consciousness, but to determine how and when it may appear, and leverage it for collective benefit.

Rather than building machines that appear seamlessly and consistently conscious, HCI could focus on revealing the multifaceted dimensions of consciousness, and so exposing the 'seams' of machine consciousness [15]. While some may see the decomposition of traits that trigger PMC as threatening the mystical or ambiguous essence of consciousness, it could also afford greater transparency for users. Machines that appear as conscious may be accompanied by a list of components that make them appear so — 'Alexa Edition 34.5 includes features of advanced search and explanation functions, emotional recognition and emulation, independence of

music preferences (techno-jazz), and spontaneous contextualised witticisms.’

At a collective level, HCI research could inform the development of regulation and policy surrounding PMC. Perhaps certain traits of consciousness will be deemed appropriate in particular contexts, such as feeling in robo-poets or android nurses, or spontaneity in machine comedians and AI improvisers. In other scenarios, traits of consciousness may be deemed inappropriate, such as independence or spontaneity in machines designed to carry out rituals or dangerous procedures. To develop and find evidence for these principles, HCI could join forces with legal experts and international regulatory bodies. Through an explicit focus on the real, observable phenomena of PMC, HCI can guide and explain the perception of consciousness in individual machines and at a collective level.

7 CONCLUSION

HCI must intellectually prepare for a future inhabited by machines that are perceived as conscious. Importantly, we make the case that the perception of machine consciousness is more critical for HCI than the ontological existence of machine consciousness. Similarly to past research achievements stimulated by the vision of Ubicomp [99], we show how studying machine consciousness right here, right now, is a means for HCI scholarship to grow and demonstrate relevance. To jump-start an HCI discourse on the perceptions of machine consciousness, we conducted a survey which showed that people already ascribe a degree of consciousness to existing technologies. Based on the results of the survey, we found that people exhibited a range of attitudes towards machine consciousness, ranging from reverence to denial. We present five dynamic tensions which offer the necessary perspective to study the confluences and contradictions inherent in the study of machine consciousness. Next, we contribute four challenges for HCI in an era of possible machine consciousness: Hierarchy and (In)equality, Empathy and Relationship, Authorship and Responsibility, and Stewardship of the Perception of Machine Consciousness. Our empirical and theoretical contributions provide a catalyst for HCI to build an important and relevant voice in the wider debate on machine consciousness.

ACKNOWLEDGMENTS

Thank you to Carolin Stellmacher and Johannes Schöning who read an earlier version of this work and provided great recommendations. Further thanks to Johannes Schöning and Yvonne Rogers who introduced the authors, and created the space for them to deliver this research as part of the Excellence Chair at University of Bremen, Germany. The first author would also like to thank Jeremy Skipper for endlessly fascinating discussions of the neuroscience of consciousness and psychedelic experiences. This research was supported by the Swedish Research Council (2022-03196) and the Leverhulme Ecological Study of the Brain Doctoral Training Programme (grant number DS-2017-026).

REFERENCES

- [1] Amal Abdulrahman, Deborah Richards, Hedieh Ranjartabar, and Samuel Mascarenhas. 2019. Belief-Based Agent Explanations to Encourage Behaviour Change. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents* (Paris, France) (IVA '19). Association for Computing Machinery, New York, NY, USA, 176–178. <https://doi.org/10.1145/3308532.3329444>
- [2] Macdonald Alex Garland, Andrew and Allan Reich. 2015. *Ex Machina*.
- [3] Judd Antin and Aaron Shaw. 2012. Social Desirability Bias and Self-Reports of Motivation: A Study of Amazon Mechanical Turk in the US and India. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 2925–2934. <https://doi.org/10.1145/2207676.2208699>
- [4] Adam Arico, Brian Fiala, Robert F Goldberg, and Shaun Nichols. 2011. The folk psychology of consciousness. *Mind & Language* 26, 3 (2011), 327–352. <https://doi.org/10.1111/j.1468-0017.2011.01420.x>
- [5] Neil Barton. 2020. Absence perception and the philosophy of zero. *Synthese* 197, 9 (2020), 3823–3850. <https://doi.org/10.1007/s11229-019-02220-x>
- [6] Eric P. S. Baumer, Mark Blythe, and Theresa Jean Tanenbaum. 2020. Evaluating Design Fiction: The Right Tool for the Job. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (DIS '20). Association for Computing Machinery, New York, NY, USA, 1901–1913. <https://doi.org/10.1145/3357236.3395464>
- [7] Hugh Beyer and Karen Holtzblatt. 1999. Contextual design. *interactions* 6, 1 (1999), 32–42. <https://doi.org/10.1145/291224.291229>
- [8] Jonathan Birch, Alexandra K Schnell, and Nicola S Clayton. 2020. Dimensions of animal consciousness. *Trends in cognitive sciences* 24, 10 (2020), 789–801. <https://doi.org/10.1126/science.1134475>
- [9] Ann Blandford, Dominic Furniss, and Stephann Makri. 2016. Qualitative HCI research: Going behind the scenes. *Synthesis lectures on human-centered informatics* 9, 1 (2016), 1–115.
- [10] Stefan Blomkvist. 2002. Persona—an overview. Retrieved November 22 (2002), 2004. <https://www.it.uu.se/edu/course/homepage/hcidist/vt05/Persona-overview.pdf>
- [11] Nick Bostrom. 2017. *Superintelligence*. Dunod.
- [12] The Editors of Encyclopaedia Britannica. 2022. "technology".
- [13] Alexander Campolo and Kate Crawford. 2020. Enchanted determinism: Power without responsibility in artificial intelligence. *Engaging Science, Technology, and Society* 6 (2020), 1–19. <https://doi.org/10.17351/ests2020.277>
- [14] David J Chalmers. 1995. Facing up to the problem of consciousness. *Journal of consciousness studies* 2, 3 (1995), 200–219.
- [15] Matthew Chalmers. 2003. Seamless design and ubicomp infrastructure. In *Proceedings of Ubicomp 2003 workshop at the crossroads: The interaction of HCI and systems issues in Ubicomp*. Citeseer, 577–584. <https://doi.org/10.1049/ic:20030140>
- [16] Zhifa Chen, Yichen Lu, Mika P. Nieminen, and Andrés Lucero. 2020. Creating a Chatbot for and with Migrants: Chatbot Personality Drives Co-Design Activities. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference* (Eindhoven, Netherlands) (DIS '20). Association for Computing Machinery, New York, NY, USA, 219–230. <https://doi.org/10.1145/3357236.3395495>
- [17] Eugene Ch'ng. 2019. Art by Computing Machinery: Is Machine Art Acceptable in the Artworld? *ACM Trans. Multimedia Comput. Commun. Appl.* 15, 2s, Article 59 (jul 2019), 17 pages. <https://doi.org/10.1145/3326338>
- [18] Grant Clauser. 2019. Amazon's Alexa Never Stops Listening to You. Should You Worry? *Wirecutter: The New York Times* (2019). <https://www.nytimes.com/wirecutter/blog/amazons-alexa-never-stops-listening-to-you/>
- [19] J. Comans, M. M van Paassen, and M. Mulder. 2010. Pilot Workload Monitoring and Adaptive Aviation Automation: A Solution Space-Based Approach. In *Proceedings of the 28th Annual European Conference on Cognitive Ergonomics* (Delft, Netherlands) (ECCE '10). Association for Computing Machinery, New York, NY, USA, 245–250. <https://doi.org/10.1145/1962300.1962351>
- [20] Kenneth James Williams Craik. 1967. *The nature of explanation*. Vol. 445. CUP Archive.
- [21] Francis Crick and Christof Koch. 2003. A framework for consciousness. *Nature neuroscience* 6, 2 (2003), 119–126. <https://doi.org/10.1038/nn0203-119>
- [22] Matthew Crosby. 2019. Why Artificial Consciousness Matters.. In *AAAI Spring Symposium: Towards Conscious AI Systems*.
- [23] James L. Crowley, Joëlle Coutaz, and François Bérard. 2000. Perceptual User Interfaces: Things That See. *Commun. ACM* 43, 3 (mar 2000), 54–ff. <https://doi.org/10.1145/330534.330540>
- [24] Felipe P. da Silva, Alan Oliveria de Sá, Nadia Nedjah, and Luiza de Macedo Mourelle. 2014. An Efficient Parallel Yet Pipelined Reconfigurable Architecture for M-PLN Weightless Neural Networks. In *Proceedings of the 27th Symposium on Integrated Circuits and Systems Design* (Aracaju, Brazil) (SBCCI '14). Association for Computing Machinery, New York, NY, USA, Article 30, 7 pages. <https://doi.org/10.1145/2660540.2660998>
- [25] Nour Dabbour and Anil Somayaji. 2020. Towards In-Band Non-Cryptographic Authentication. In *New Security Paradigms Workshop 2020* (Online, USA) (NSPW '20). Association for Computing Machinery, New York, NY, USA, 20–33. <https://doi.org/10.1145/3442167.3442180>
- [26] Daniel Clement Dennett. 1987. *The intentional stance*. MIT press.
- [27] Andrew Dillon and Charles Watson. 1996. User analysis in HCI—the historical lessons from individual differences research. *International journal of human-computer studies* 45, 6 (1996), 619–637. <https://doi.org/10.1006/ijhc.1996.0071>
- [28] Steve Donaldson, Jesse Kawell, and Chris Walling. 2011. Computer, Know Thyself: Exploring Consciousness via Self-Aware Machines. In *Proceedings of*

- the 49th Annual Southeast Regional Conference (Kennesaw, Georgia) (ACM-SE '11). Association for Computing Machinery, New York, NY, USA, 70–74. <https://doi.org/10.1145/2016039.2016064>
- [29] Judith Dörrenbächer, Diana Löffler, and Marc Hassenzahl. 2020. Becoming a Robot - Overcoming Anthropomorphism with Techno-Mimesis. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376507>
- [30] Mike Elgan. 2018. The case against teaching kids to be polite to Alexa. *Fast Company* (2018). <https://www.fastcompany.com/40588020/the-case-against-teaching-kids-to-be-polite-to-alexa>
- [31] Eric Elliot. 2020. What It's Like To be a Computer: An Interview with GPT-3. https://www.youtube.com/watch?v=PqbB07n_uQ4
- [32] Chris Elsdén, David Chatting, Abigail C. Durrant, Andrew Garbett, Bettina Nissen, John Vines, and David S. Kirk. 2017. On Speculative Enactments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New York, NY, USA, 5386–5399. <https://doi.org/10.1145/3025453.3025503>
- [33] Nicholas Epley, Adam Waytz, and John T. Cacioppo. 2007. On seeing human: a three-factor theory of anthropomorphism. *Psychological review* 114, 4 (2007), 864. <https://doi.org/10.1037/0033-295X.114.4.864>
- [34] Friederike Eyssel, Dieter Kuchenbrandt, and Simon Bobinger. 2011. Effects of Anticipated Human-Robot Interaction and Predictability of Robot Behavior on Perceptions of Anthropomorphism. In *Proceedings of the 6th International Conference on Human-Robot Interaction* (Lausanne, Switzerland) (HRI '11). Association for Computing Machinery, New York, NY, USA, 61–68. <https://doi.org/10.1145/1957656.1957673>
- [35] Valentina Franzoni, Alfredo Milani, and Jordi Vallverdú. 2017. Emotional Affordances in Human-Machine Interactive Planning and Negotiation. In *Proceedings of the International Conference on Web Intelligence* (Leipzig, Germany) (WI '17). Association for Computing Machinery, New York, NY, USA, 924–930. <https://doi.org/10.1145/3106426.3109421>
- [36] Sophie Freeman, Martin Gibbs, and Björn Nansen. 2022. 'Don't mess with my algorithm': Exploring the relationship between listeners and automated curation and recommendation on music streaming services. *First Monday* (2022). <https://doi.org/10.5210/fm.v27i1.11783>
- [37] Anne Gerdes. 2016. The Issue of Moral Consideration in Robot Ethics. *SIGCAS Comput. Soc.* 45, 3 (jan 2016), 274–279. <https://doi.org/10.1145/2874239.2874278>
- [38] Heather M. Gray, Kurt Gray, and Daniel M. Wegner. 2007. Dimensions of mind perception. *science* 315, 5812 (2007), 619–619. <https://doi.org/10.1126/science.1134475>
- [39] Michael SA Graziano. 2019. *Rethinking consciousness: a scientific theory of subjective experience*. WW Norton & Company.
- [40] Peter H. Greene. 1959. An Approach to Computers That Perceive, Learn, and Reason. In *Papers Presented at the March 3-5, 1959, Western Joint Computer Conference* (San Francisco, California) (IRE-AIEE-ACM '59 (Western)). Association for Computing Machinery, New York, NY, USA, 181–186. <https://doi.org/10.1145/1457838.1457870>
- [41] Stewart Elliott Guthrie. 1997. Anthropomorphism: A definition and a theory. (1997).
- [42] Nick Haslam, Stephen Loughnan, Yoshihisa Kashima, and Paul Bain. 2008. Attributing and denying humanness to others. *European review of social psychology* 19, 1 (2008), 55–85. <https://doi.org/10.1080/10463280801981645>
- [43] Marc Hassenzahl. 2003. The thing and I: understanding the relationship between user and product. In *Funology*. Springer, 31–42. https://doi.org/10.1007/978-3-319-68213-6_19
- [44] David Hauser, Gabriele Paolacci, and Jesse Chandler. 2019. Common concerns with MTurk as a participant pool: Evidence and solutions. In *Handbook of research methods in consumer psychology*. Routledge, 319–337. <https://doi.org/10.4324/9781351137713-17>
- [45] Virginia Heffernan. 2020. My Roomba Has Achieved Enlightenment. *WIRED* (2020). <https://www.wired.com/story/roomba-robot-consciousness-enlightenment/>
- [46] Frank Hegel, Soeren Krach, Tilo Kircher, Britta Wrede, and Gerhard Sagerer. 2008. Theory of Mind (ToM) on Robots: A Functional Neuroimaging Study. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction* (Amsterdam, The Netherlands) (HRI '08). Association for Computing Machinery, New York, NY, USA, 335–342. <https://doi.org/10.1145/1349822.1349866>
- [47] Fritz Heider and Marianne Simmel. 1944. An experimental study of apparent behavior. *The American journal of psychology* 57, 2 (1944), 243–259. <https://doi.org/10.2307/1416950>
- [48] Daniel T. Heinze, Alexander Turchin, and V. Jagannathan. 2006. Automated Interpretation of Clinical Encounters with Cultural Cues and Electronic Health Record Generation. In *Proceedings of the Workshop on Medical Speech Translation* (New York, New York) (MST '06). Association for Computational Linguistics, USA, 20–27.
- [49] Laura M. Hiatt and J. Gregory Trafton. 2015. Understanding Second-Order Theory of Mind. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts* (Portland, Oregon, USA) (HRI'15 Extended Abstracts). Association for Computing Machinery, New York, NY, USA, 167–168. <https://doi.org/10.1145/2701973.2702030>
- [50] Robin K. Hill and Carlos Baquero. 2021. Seeking out Camille, and Being Open to Others. *Commun. ACM* 64, 12 (nov 2021), 14–15. <https://doi.org/10.1145/3490151>
- [51] Michael Hind. 2019. Explaining Explainable AI. *XRDS* 25, 3 (April 2019), 16–19. <https://doi.org/10.1145/3313096>
- [52] Bryce Huebner. 2010. Commonsense concepts of phenomenal consciousness: Does anyone care about functional zombies? *Phenomenology and the cognitive sciences* 9, 1 (2010), 133–155. <https://doi.org/10.1007/s11097-009-9126-6>
- [53] Takashi Ikegami. 2010. Studying a Self-Sustainable System by Making a Mind Time Machine. In *Workshop on Self-Sustaining Systems* (Tokyo, Japan) (S3 '10). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/1942793.1942794>
- [54] Theodore Jensen, Mohammad Maifi Hasan Khan, Md Abdullah Al Fahim, and Yusuf Albayram. 2021. Trust and Anthropomorphism in Tandem: The Interrelated Nature of Automated Agent Appearance and Reliability in Trustworthiness Perceptions. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (DIS '21). Association for Computing Machinery, New York, NY, USA, 1470–1480. <https://doi.org/10.1145/3461778.3462102>
- [55] Haiyan Jia, Mu Wu, Eunhwa Jung, Alice Shapiro, and S. Shyam Sundar. 2012. Balancing Human Agency and Object Agency: An End-User Interview Study of the Internet of Things. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (Pittsburgh, Pennsylvania) (UbiComp '12). Association for Computing Machinery, New York, NY, USA, 1185–1188. <https://doi.org/10.1145/2370216.2370470>
- [56] Brigitte Jordan and Austin Henderson. 1995. Interaction analysis: Foundations and practice. *The journal of the learning sciences* 4, 1 (1995), 39–103. https://doi.org/10.1207/s15327809jls0401_2
- [57] Filip Jurcicek, Simon Keizer, Milica Gašić, Francois Mairesse, Blaise Thomson, Kai Yu, and Steve Young. 2011. Real user evaluation of spoken dialogue systems using Amazon Mechanical Turk. In *Proceedings of INTERSPEECH*, Vol. 11.
- [58] Evangelos Karapanos, John Zimmerman, Jodi Forlizzi, and Jean-Bernard Martens. 2009. User Experience over Time: An Initial Framework. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) (CHI '09). Association for Computing Machinery, New York, NY, USA, 729–738. <https://doi.org/10.1145/1518701.1518814>
- [59] Junhan Kim, Yoojung Kim, Byungjoon Kim, Sukyoung Yun, Minjoon Kim, and Joongseok Lee. 2018. Can a Machine Tend to Teenagers' Emotional Needs? A Study with Conversational Agents. 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, 1–6. <https://doi.org/10.1145/3170427.3188548>
- [60] Kyung-Joong Kim and Hod Lipson. 2009. Towards a Simple Robotic Theory of Mind. In *Proceedings of the 9th Workshop on Performance Metrics for Intelligent Systems* (Gaithersburg, Maryland) (PerMIS '09). Association for Computing Machinery, New York, NY, USA, 131–138. <https://doi.org/10.1145/1865909.1865937>
- [61] Elsa A. Kirchner, Stephen H. Fairclough, and Frank Kirchner. 2019. *Embedded Multimodal Interfaces in Robotics: Applications, Future Trends, and Societal Implications*. Association for Computing Machinery and Morgan & Claypool, 523–576. <https://doi.org/10.1145/3233795.3233810>
- [62] C. Kline. 1987. Supercomputers on the Internet: A Case Study. In *Proceedings of the ACM Workshop on Frontiers in Computer Communications Technology* (Stowe, Vermont, USA) (SIGCOMM '87). Association for Computing Machinery, New York, NY, USA, 27–33. <https://doi.org/10.1145/55482.55487>
- [63] Ralf Klischewski. 2019. Will the Government Machine Turn into a Monster?. In *Proceedings of the 20th Annual International Conference on Digital Government Research* (Dubai, United Arab Emirates) (dg.o 2019). Association for Computing Machinery, New York, NY, USA, 306–313. <https://doi.org/10.1145/3325112.3325221>
- [64] Christof Koch. 2019. *The feeling of life itself: why consciousness is widespread but can't be computed*. Mit Press.
- [65] Marion Koelle, Swamy Ananthanarayan, and Susanne Boll. 2020. Social Acceptability in HCI: A Survey of Methods, Measures, and Design Strategies. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–19. <https://doi.org/10.1145/3313831.3376162>
- [66] Christos Kouroupetroglou, Michail Salamapasis, and Athanasios Manitsaris. 2006. A Semantic-Web Based Framework for Developing Applications to Improve Accessibility in the WWW. In *Proceedings of the 2006 International Cross-Disciplinary Workshop on Web Accessibility (W4A): Building the Mobile Web: Rediscovering Accessibility?* (Edinburgh, United Kingdom) (W4A '06). Association for Computing Machinery, New York, NY, USA, 98–108. <https://doi.org/10.1145/1133219.1133238>
- [67] Leantos Kyriakoullis and Panayiotis Zaphiris. 2016. Culture and HCI: a review of recent cultural studies in HCI and social networks. *Universal Access in the*

- Information Society 15, 4 (2016), 629–642. <https://doi.org/10.1007/s10209-015-0445-9>
- [68] Guy Laban. 2021. Perceptions of Anthropomorphism in a Chatbot Dialogue: The Role of Animacy and Intelligence. In *Proceedings of the 9th International Conference on Human-Agent Interaction* (Virtual Event, Japan) (HAI '21). Association for Computing Machinery, New York, NY, USA, 305–310. <https://doi.org/10.1145/3472307.3484686>
- [69] Carine Lallemand and Emeline Mercier. 2022. Optimizing the Use of the Sentence Completion Survey Technique in User Research: A Case Study on the Experience of E-Reading. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 319, 18 pages. <https://doi.org/10.1145/3491102.3517718>
- [70] Effie Lai-Chong Law, Virpi Roto, Marc Hassenzahl, Arnold POS Vermeeren, and Joke Kort. 2009. Understanding, scoping and defining user experience: a survey approach. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 719–728. <https://doi.org/10.1145/1518701.1518813>
- [71] Hong Lu, Weibing Hu, and Chaochao Pan. 2021. Artificial Intelligence and Its Self-Consciousness. In *2021 Workshop on Algorithm and Big Data* (Fuzhou, China) (WABD 2021). Association for Computing Machinery, New York, NY, USA, 67–69. <https://doi.org/10.1145/3456389.3456402>
- [72] Andrés Lucero. 2015. Using affinity diagrams to evaluate interactive prototypes. In *ITIP conference on human-computer interaction*. Springer, 231–248. https://doi.org/10.1007/978-3-319-22668-2_19
- [73] Karl F MacDorman and Hiroshi Ishiguro. 2006. The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies* 7, 3 (2006), 297–337. <https://doi.org/10.1075/is.7.3.03mac>
- [74] Clara Mancini, Yvonne Rogers, Arosha K Bandara, Tony Coe, Lukasz Jedrzejczyk, Adam N Joinson, Blaine A Price, Keerthi Thomas, and Bashar Nuseibeh. 2010. Contrivance: exploring users' reactions to futuristic technology. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 153–162. <https://doi.org/10.1145/1753326.1753350>
- [75] Gary Marcus, Ernest Davis, and Scott Aaronson. 2022. A very preliminary analysis of DALL-E 2. <https://doi.org/10.48550/ARXIV.2204.13807>
- [76] Alice Martin, Mathieu Magnaudet, and Stephane Conversy. 2021. Vers La Complétude Interactive: Exigences Pour Une Machine Abstraite Orientée Interaction: Toward Interactive Completeness: Requirements for an Interactive Abstract Machine. In *32e Conférence Francophone Sur l'Interaction Homme-Machine* (Virtual Event, France) (IHM '21). Association for Computing Machinery, New York, NY, USA, Article 9, 6 pages. <https://doi.org/10.1145/3451148.3458644>
- [77] Justin Matejka, Michael Glueck, Toví Grossman, and George Fitzmaurice. 2016. The Effect of Visual Appearance on the Performance of Continuous Sliders and Visual Analogue Scales. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 5421–5432. <https://doi.org/10.1145/2858036.2858063>
- [78] Paul K McClure. 2018. "You're fired," says the robot: The rise of automation in the workplace, technophobes, and fears of unemployment. *Social Science Computer Review* 36, 2 (2018), 139–156. <https://doi.org/10.1177/0894439317698637>
- [79] Daniel McDuff and Mary Czerwinski. 2018. Designing Emotionally Sentient Agents. *Commun. ACM* 61, 12 (nov 2018), 74–83. <https://doi.org/10.1145/3186591>
- [80] Elisa D Mekler and Kasper Hornbæk. 2016. Momentary pleasure or lasting meaning? Distinguishing eudaimonic and hedonic user experiences. In *Proceedings of the 2016 chi conference on human factors in computing systems*. 4509–4520. <https://doi.org/10.1145/2858036.2858225>
- [81] Merriam-Webster. [n.d.]. *Sentience*. <https://www.merriam-webster.com/dictionary/sentience>
- [82] Cade Metz. 2022. A.I. Is Not Sentient. Why Do People Say It Is? (August 2022). <https://www.nytimes.com/2022/08/05/technology/ai-sentient-google.html>
- [83] Selina Meyer, David Elswiler, Bernd Ludwig, Marcos Fernandez-Pichel, and David E. Losada. 2022. Do We Still Need Human Assessors? Prompt-Based GPT-3 User Simulation in Conversational AI. In *Proceedings of the 4th Conference on Conversational User Interfaces* (Glasgow, United Kingdom) (CUI '22). Association for Computing Machinery, New York, NY, USA, Article 8, 6 pages. <https://doi.org/10.1145/3543829.3544529>
- [84] Albert Michotte. 2017. *The perception of causality*. Routledge. <https://doi.org/10.4324/9781315519050>
- [85] Thomas Nagel. 1974. What is it like to be a bat? *The philosophical review* 83, 4 (1974), 435–450. <https://doi.org/10.2307/2183914>
- [86] Jakob Nielsen. 2005. Ten usability heuristics.
- [87] Jasmin Niess and Paweł Wóźniak. 2020. Embracing companion technologies. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–11. <https://doi.org/10.1145/3419249.3420134>
- [88] Meghan O'Gieblyn. 2021. *God Human Animal Machine. Technology, Metaphor, and the Search for Meaning*. Penguin Random House LLC, New York.
- [89] Tetsuo Ono, Michita Imai, and Ryohei Nakatsu. 2000. Reading a robot's mind: A model of utterance understanding based on the theory of mind mechanism. *Advanced Robotics* 14, 4 (2000), 311–326. <https://doi.org/10.1163/15685300741609>
- [90] Sharon Oviatt, Björn Schuller, Philip R. Cohen, Daniel Sonntag, Gerasimos Potamianos, and Antonio Krüger (Eds.). 2019. *The Handbook of Multimodal-Multisensor Interfaces: Language Processing, Software, Commercialization, and Emerging Directions*. Association for Computing Machinery and Morgan & Claypool. <https://doi.org/10.1145/3233795>
- [91] Donna Patterson. 2004. Encouraging and Assisting Faculty Incorporation of Innovative Classroom Technologies. In *Proceedings of the 32nd Annual ACM SIGUCCS Conference on User Services* (Baltimore, MD, USA) (SIGUCCS '04). Association for Computing Machinery, New York, NY, USA, 79–81. <https://doi.org/10.1145/1027802.1027823>
- [92] A. Pawlowski. 2019. Should you be polite to your smart speakers? Why many people say 'please' to Alexa. *Today* (2019). <https://www.today.com/health/should-you-be-polite-your-smart-speakers-why-many-people-t169365>
- [93] Bogdan Popoveniuc. 2013. Pro and Cons Singularity: Kurzweil's Theory and Its Critics. In *Proceedings of the Virtual Reality International Conference: Laval Virtual* (Laval, France) (VRIC '13). Association for Computing Machinery, New York, NY, USA, Article 30, 6 pages. <https://doi.org/10.1145/2466816.2466848>
- [94] 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* (Denver, Colorado, USA) (CHI EA '17). Association for Computing Machinery, New York, NY, USA, 2853–2859. <https://doi.org/10.1145/3027063.3053246>
- [95] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-Shot Text-to-Image Generation. *CoRR* abs/2102.12092 (2021). arXiv:2102.12092 <https://arxiv.org/abs/2102.12092>
- [96] Antti Revonsuo and Matti Kamppinen. 2013. *Consciousness in philosophy and cognitive neuroscience*. Psychology Press.
- [97] Julie Rico and Stephen Brewster. 2010. Usable gestures for mobile interfaces: evaluating social acceptability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 887–896. <https://doi.org/10.1145/1753326.1753458>
- [98] Laurel D. Riek, Tal-Chen Rabinowitch, Bhismadev Chakrabarti, and Peter Robinson. 2009. How Anthropomorphism Affects Empathy toward Robots. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction* (La Jolla, California, USA) (HRI '09). Association for Computing Machinery, New York, NY, USA, 245–246. <https://doi.org/10.1145/1514095.1514158>
- [99] Yvonne Rogers. 2006. Moving on from Weiser's Vision of Calm Computing: engaging UbiComp experiences. *UbiComp '06* (2006), 404–421. https://doi.org/10.1007/11853565_24 ISBN: 978-3-540-39634-5.
- [100] Beat Rossmay, Sarah Theres Völkel, Elias Naphausen, Patricia Kimm, Alexander Wiethoff, and Andreas Muxel. 2020. Punishable AI: Examining Users' Attitude Towards Robot Punishment. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 179–191. <https://doi.org/10.1145/3357236.3395542>
- [101] Marco C. Rozendaal, Boudewijn Boon, and Victor Kaptein. 2019. Objects with Intent: Designing Everyday Things as Collaborative Partners. *ACM Trans. Comput.-Hum. Interact.* 26, 4, Article 26 (jun 2019), 33 pages. <https://doi.org/10.1145/3325277>
- [102] Cynthia Rudin. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1, 5 (2019), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [103] Colin T. Schmidt. 2020. Extensions of Mind: Relating to Data and Others. In *Proceedings of the 11th Augmented Human International Conference* (Winnipeg, Manitoba, Canada) (AH '20). Association for Computing Machinery, New York, NY, USA, Article 13, 7 pages. <https://doi.org/10.1145/3396339.3396402>
- [104] Brian J Scholl and Patrice D Tremoulet. 2000. Perceptual causality and animacy. *Trends in cognitive sciences* 4, 8 (2000), 299–309. [https://doi.org/10.1016/S1364-6613\(00\)01506-0](https://doi.org/10.1016/S1364-6613(00)01506-0)
- [105] Amazon Science. 2020. *Alexa Prize Socialbot Grand Challenge 3 Finals*. <https://www.youtube.com/watch?v=Rh3D756garE>
- [106] Anil Seth. 2021. *Being you: A new science of consciousness*. Penguin. <https://doi.org/10.1353/wlt.2022.0093>
- [107] Anil K Seth, Eugene Izhikevich, George N Reeke, and Gerald M Edelman. 2006. Theories and measures of consciousness: an extended framework. *Proceedings of the National Academy of Sciences* 103, 28 (2006), 10799–10804. <https://doi.org/10.1073/pnas.0604347103>
- [108] William Seymour and Max Van Kleek. 2020. Does Siri Have a Soul? Exploring Voice Assistants Through Shinto Design Fictions. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3334480.3381809>
- [109] 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 (oct 2021), 16 pages. <https://doi.org/10.1145/3479515>
- [110] Hanieh Shakeri, Carman Neustaedter, and Steve DiPaola. 2021. SAGA: Collaborative Storytelling with GPT-3. In *Companion Publication of the 2021 Conference*

- on Computer Supported Cooperative Work and Social Computing (Virtual Event, USA) (CSCW '21). Association for Computing Machinery, New York, NY, USA, 163–166. <https://doi.org/10.1145/3462204.3481771>
- [111] Janelle Shane. 2022. Interview with a squirrel. (June 2022). <https://www.aiweirdness.com/interview-with-a-squirrel/>
- [112] Helen Sharp, Jennifer Preece, and Yvonne Rogers. 2019. *Interaction Design: Beyond Human - Computer Interaction* (5th ed.). Wiley Publishing.
- [113] Sara J Shettleworth. 2010. Clever animals and killjoy explanations in comparative psychology. *Trends in cognitive sciences* 14, 11 (2010), 477–481. <https://doi.org/10.1016/j.tics.2010.07.002>
- [114] Anastasia Siapka. 2022. Towards a Feminist Metaethics of AI. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (AIES '22). Association for Computing Machinery, New York, NY, USA, 665–674. <https://doi.org/10.1145/3514094.3534197>
- [115] Michelle Sleed, Kevin Durrheim, Anita Kriel, Vernon Solomon, and Veronica Baxter. 2002. The effectiveness of the vignette methodology: A comparison of written and video vignettes in eliciting responses about date rape. *South African Journal of Psychology* 32, 3 (2002), 21–28. <https://doi.org/10.10520/EJC98194>
- [116] Vincent Landay Spike Jonze, Megan Ellison. 2013. Her.
- [117] Daniel Steinbock and Sridhar Rao. 2020. User Experience Design Patterns for Pseudo-Sentient Agents. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3334480.3381061>
- [118] Mengmeng Sun, Chenyao Tian, Liyang Mao, Xianghe Meng, Xingjian Shen, Bo Hao, Xin Wang, Hui Xie, and Li Zhang. 2022. Reconfigurable Magnetic Slime Robot: Deformation, Adaptability, and Multifunction. *Advanced Functional Materials* (2022), 2112508. <https://doi.org/10.1002/adfm.202112508>
- [119] Justin Sytsma. 2010. Dennett's theory of the folk theory of consciousness. *Journal of Consciousness Studies* 17, 3-4 (2010), 107–130.
- [120] Amelia Tait. 2022. 'I am, in fact, a person': can artificial intelligence ever be sentient? (2022). <https://www.theguardian.com/technology/2022/aug/14/can-artificial-intelligence-ever-be-sentient-googles-new-ai-program-is-raising-questions>
- [121] Leila Takayama. 2012. Perspectives on agency interacting with and through personal robots. In *Human-computer interaction: the agency perspective*. Springer, 195–214. https://doi.org/10.1007/978-3-642-25691-2_8
- [122] Allie Terry-Fritsch. 2017. Animal Trials, Humiliation Rituals, and the Sensuous Suffering of Criminal Offenders in Late Medieval and Early Modern Europe. In *Visualizing Sensuous Suffering and Affective Pain in Early Modern Europe and the Spanish Americas*. Brill, 53–81. https://doi.org/10.1163/9789004360686_004
- [123] Wenxiu Tian, Hongxin Wang, and Yunfeng Bao. 2021. College Physics Experiment Teaching in Application-Oriented College. In *The Sixth International Conference on Information Management and Technology* (Jakarta, Indonesia) (ICIMTECH 21). Association for Computing Machinery, New York, NY, USA, Article 277, 5 pages. <https://doi.org/10.1145/3465631.3465949>
- [124] Nitasha Tiku. 2022. The Google engineer who thinks the company's AI has come to life. (June 2022). <https://www.washingtonpost.com/technology/2022/06/11/google-ai-lambda-blake-lemoine/>
- [125] Giulio Tononi. 2015. Integrated information theory. *Scholarpedia* 10, 1 (2015), 4164. <https://doi.org/10.1038/nrn.2016.44>
- [126] Videvo. [n.d.]. *Close Up Shot of Robotic Vacuum Cleaner Cleaning a Carpet*. <https://www.videvo.net/video/close-up-shot-of-robotic-vacuum-cleaner-cleaning-a-carpet/993158/>
- [127] Google Books Ngram Viewer. [n.d.]. *Technology V Machine Word Count Comparison, 1950-2019*. https://books.google.com/ngrams/graph?content=technology%2Cmachine&year_start=1950&year_end=2019&case_insensitive=on&corpus=26&smoothing=3 Date accessed: 05/12/2022.
- [128] Anna-Lisa Vollmer. 2018. Fears of Intelligent Robots. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (Chicago, IL, USA) (HRI '18). Association for Computing Machinery, New York, NY, USA, 273–274. <https://doi.org/10.1145/3173386.3177067>
- [129] Rina R. Wehbe, Edward Lank, and Lennart E. Nacke. 2017. Left Them 4 Dead: Perception of Humans versus Non-Player Character Teammates in Cooperative Gameplay. In *Proceedings of the 2017 Conference on Designing Interactive Systems* (Edinburgh, United Kingdom) (DIS '17). Association for Computing Machinery, New York, NY, USA, 403–415. <https://doi.org/10.1145/3064663.3064712>
- [130] Wei Wei, Jiayi Liu, Xianling Mao, Guibing Guo, Feida Zhu, Pan Zhou, and Yuchong Hu. 2019. Emotion-Aware Chat Machine: Automatic Emotional Response Generation for Human-like Emotional Interaction. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (CIKM '19). Association for Computing Machinery, New York, NY, USA, 1401–1410. <https://doi.org/10.1145/3357384.3357937>
- [131] Juyang (John) Weng. 2022. An Algorithmic Theory for Conscious Learning. In *2022 The 3rd International Conference on Artificial Intelligence in Electronics Engineering* (Bangkok, Thailand) (AIEE 2022). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3512826.3512827>
- [132] Jiri Wiedermann. 1999. Simulating the Mind: A Gauntlet Thrown to Computer Science. *ACM Comput. Surv.* 31, 3es (sep 1999), 16–es. <https://doi.org/10.1145/333580.333595>
- [133] Eva Wiseman. 2021. Is robot therapy the future? *The Guardian* (August 2021). www.theguardian.com/society/2021/aug/08/is-robot-therapy-the-future
- [134] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, BC, Canada) (CHI '11). Association for Computing Machinery, New York, NY, USA, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [135] Richmond Y Wong and Vera Khovanskaya. 2018. Speculative design in HCI: from corporate imaginations to critical orientations. In *New Directions in Third Wave Human-Computer Interaction: Volume 2-Methodologies*. Springer, 175–202. https://doi.org/10.1007/978-3-319-73374-6_10
- [136] Paweł W Woźniak, Jakob Karolus, Florian Lang, Caroline Eckerth, Johannes Schöning, Yvonne Rogers, and Jasmin Niess. 2021. Creepy technology: what is it and how do you measure it?. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3411764.3445299>
- [137] Nur Yildirim, James McCann, and John Zimmerman. 2020. Digital Fabrication Tools at Work: Probing Professionals' Current Needs and Desired Futures. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376621>
- [138] Jakub Zlotowski, Ewald Strasser, and Christoph Bartneck. 2014. Dimensions of anthropomorphism: from humanness to humanlikeness. In *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 66–73. <https://doi.org/10.1145/2559636.2559679>