

# Should You Trust Your Voice Assistant? It’s Complicated, but No

Filippos Stamatiou<sup>1</sup> and Xenofon Karakonstantis<sup>2</sup>

**Abstract**—The widespread use of voice-assisted applications using artificial intelligence raises questions about the dynamics of trust and reliance on these systems. While users often rely on these applications for help, instances where users face unforeseen risks and heightened challenges have sparked conversations about the importance of fostering trustworthy artificial intelligence. In this paper, we argue that the prevailing narrative of trust and trustworthiness in relation to artificial intelligence, particularly voice assistants, is misconstrued and fundamentally misplaced. Drawing on insights from philosophy and artificial intelligence literature, we contend that artificial intelligence systems do not meet the criteria for participating in a relationship of trust with human users. Instead, a narrative of reliance is more appropriate. However, we investigate the matter further, to explore why the trust/trustworthiness narrative persists, focusing on the unique social dynamics of interactions with voice assistants. We identify factors such as diverse modalities and complexity, social aspects of voice assistants, and issues of uncertainty, assertiveness, and transparency as contributors to the trust narrative. By disentangling these factors, we shed light on the complexities of human-computer interactions and offer insights into the implications for our relationship with artificial intelligence. We advocate for a nuanced understanding of trust and reliance in artificial intelligence systems and provide suggestions for addressing the challenges posed by the dominance of the trust/trustworthiness narrative.

## I. INTRODUCTION

Suppose that you fly into a city for the first time. You pick up a rental car and start driving towards your hotel. As the city is unfamiliar to you, you enter the hotel address in the navigation app on your phone and a number of potential routes appear for you to choose. You choose the fastest route, even though it does not follow the main highway into the city. After all, the app has live information on traffic conditions, accidents on the road, or other incidents. The fastest route is the best route, and you trust the app, or at least you think you do. You start driving while the app’s voice assistant is giving you directions: “Turn left”, “In 300 metres, take the first exit” “Continue straight”. Unbeknownst to you, the fastest route takes you right through the city’s most unsafe neighbourhood, where car hijackings are extremely common and where criminals prey on unsuspecting tourists in rental cars, just like you. Fortunately, nothing happens and you arrive safely at your hotel.

But others are less lucky, and they blame the app providers for their misfortunes. In fact, according to a recent news piece, a couple from the US is suing Google because Google Maps led them into a dangerous neighbourhood, where their

rental car was attacked and they were seriously injured and robbed of their belongings [1].

What makes users willing to take legal action against a technology company and its navigation app for leading them into an unsafe neighbourhood? Crime is unpredictable, and navigation apps do not claim to protect against hijackers. Yet, users are led to reasonably believe that they can -or even should- trust their navigation apps in meaningful ways. In this paper, we argue that while this trust is misconstrued and fundamentally misplaced, the reasons for the proliferation of the trust/trustworthiness narrative in relation to artificial intelligence (AI) boil down to how many of these apps are programmed to interact with users. We focus on voice assistants, which are a paradigmatic case of human-machine social interaction and highlight the causes, dangers, and implications of the problematic ambiguity of trust in AI.

We begin by introducing the concepts of trust and trustworthiness in the philosophical literature, as well as selected insights from the growing literature on trust in AI. We conclude that the professed demand for trustworthy AI is problematic, as machines do not meet the criteria for participating in a relationship of trust with a human user. A more appropriate attitude towards AI would be reliance.

Most work on trust and trustworthiness in AI stops here, calling for refocusing the narrative from trust and trustworthiness to reliance. However, we advance the debate by employing the case of voice assistants to explain *why* the trust/trustworthiness narrative has become dominant. We arrive at a set of factors that enable the perception of a relationship of trust between user and voice assistant. Finally, we argue that this has negative implications for our relationship with AI, now and in the future.

## II. TRUST AND TRUSTWORTHINESS

### A. Philosophy of Trust and Trustworthiness

Philosophers agree on very little when it comes to trust. However, many converge on the view that trust is an attitude that allows us to depend on others [2] and that it involves risk, usually in the form of vulnerability of the trustor towards the trustee [3]. In contrast, trustworthiness is commonly thought of as a property that a trustee possesses. In this section, we briefly review the literature on trust and trustworthiness in philosophy, both in general and specifically in relation to AI. These insights provide a first approximation of the nature of trust and the conditions under which someone -or something- may be considered trustworthy.

Richard Holton [4] argues that an attitude of trust is unique because it involves taking the *participant stance*

<sup>1</sup>Centre for the Philosophy of Artificial Intelligence, University of Copenhagen and Department of Philosophy, Stellenbosch University

<sup>2</sup>Department of Electrical and Photonics Engineering, Technical University of Denmark



uitous in the last decade. Voice assistants, including Amazon Alexa, Microsoft Cortana, Google Assistant, and Apple Siri, facilitate various tasks such as information retrieval, meeting scheduling, and hands-free calling from vehicles or homes. Employing Natural Language User Interfaces (NLUI), these assistants engage users and provide services spanning weather updates, navigation guidance, schedule management, and phone call facilitation. Figure 1 illustrates an example of a voice assistant pipeline, which includes language input through speech or text. This process often incorporates voice detection (wake word) models for activation, [26] automatic speech recognition models for converting speech to text, [27] intent recognition models for interpreting user inputs, [28] and text-to-speech models for auditory output. [29] Additionally, cloud computing is commonly leveraged for efficiency, data storage, and edge device utilization, with encryption and supplementary privacy features typically integrated for enhanced security.

In what follows, we touch upon three characteristics of voice assistants that are responsible for both their remarkable success and popularity as well as the idea that they can be trustworthy and humans should be able to trust them. Although these characteristics have already been discussed in relative detail, we suggest how they contribute to the narrative of trust and trustworthiness. Our aim is to disentangle the trust/reliance debate from the complications of mere terminology. We think the demand for trustworthy AI is not just a terminological mistake, but rather rooted in aspects of the technology and its interaction with humans, such as the ones we suggest below.

#### A. Diverse Modalities and Complexity

Voice assistants, exemplified by Amazon Alexa and Google Assistant, offer a wide range of capabilities that range from mundane tasks such as setting reminders to more complex functions such as offering medical advice. However, this breadth of functionality introduces a level of uncertainty regarding the reliability of each task. For instance, in challenging acoustic environments characterized by high levels of noise, the speech recognition component of these assistants may falter, leading to misinterpretations of user commands or requests. This challenge relates predominantly to the reliability of the system rather than an issue of trust. It is imperative to perform a granular evaluation of each component individually, including the hardware and software responsible for capturing speech signals. This approach ensures a comprehensive understanding of the system's performance, avoiding reliance solely on instances of misinterpretations by the voice assistant. Additionally, as emphasized by Kim et al. [30], while voice assistants initially captivate users with their intuitive speech-based interfaces, prolonged engagement often reveals difficulties in articulating commands clearly for users, alongside apprehensions regarding privacy and security.

Recent literature has shed light on critical sociotechnical "trust" challenges confronting audiovisual assistive technologies, particularly evident in the domain of hearing aids.

[31] These challenges encompass complexities in technology integration, cost constraints, limitations in battery life, and addressing user non-compliance. Similar challenges are encountered in the realm of voice assistants, where the reliability of each aforementioned component is paramount. However, it is unwarranted to attribute untrustworthiness to these technologies solely based on occasional performance limitations. Instead, such limitations underscore the importance of evaluating the reliability of AI assistants across a spectrum of situational contexts. Just as we recognize the multifaceted nature of hearing aids and refrain from questioning their reliability in challenging scenarios, a nuanced understanding of AI assistant technology prompts us to assess their reliability across diverse scenarios, particularly considering their operational capabilities and constraints, i.e. challenges in speech processing, natural language processing, speech synthesis, etc. Therefore, which of these components should we trust?

It is imperative to recognize that characterizing these challenges solely as matters of trust can obscure the crucial distinction between reliability and trustworthiness. This distinction is fundamental for comprehending the factors that shape technology adoption and acceptance within society. A tool is as good as each of its individual parts. Such tools are often also just one component within a broader ecosystem of technologies that aim to improve user experiences in different industries. For example, AI voice assistants are integrated in a wide range of wildly different tasks, from building design tools [32] to helping with the cognitive effort of high-risk decision making [33].

Voice assistants are a paradigmatic case of a highly complex system that performs a wide range of functions, all the while using a simple user interface consisting of just voice inputs and outputs. From the perspective of the human user, a voice assistant can help with diverse everyday tasks, from route navigation to playing music to answering factual questions. And while the inner workings of the system are obscure to the user -and to a certain degree even to the developer-, their usefulness is indisputable.

This combination of complexity and convenience is conducive to a narrative of trust and trustworthiness around voice assistants. While they do not fulfill the conditions for even a deflated notion of trust, the diversity and parallel nature of their functions create an expectation of trustworthiness in the human user. One is hard pressed to pinpoint exactly what a voice assistant is *relying on* to do. Rather, *trusting* a voice assistant seems more appropriate, even if the notion is misconstrued.

#### B. Social Aspects of Voice Assistants

A second aspect of voice assistants that gives rise to a narrative of trust and trustworthiness is their social dimension. That has already been the focus of much work (see for instance [34] for a proposed research agenda). One can ask questions, request information, complete various tasks, receive instructions or navigation, or even have an open-ended conversation with a voice assistant. Importantly, voice

assistants are purposefully designed to promote and maintain intimacy with the user [35]. Research has shown that we are motivated to use voice assistants at least partly by the social benefits they provide [36]. Whether these benefits are real and sustainable or not is debatable [35], but the fact remains that for most people, voice assistants present the first and main opportunity to interact with AI on a daily basis [37].

However, little attention has been paid to the dangers of voice assistants as social interaction partners (see, for instance, [38]). Most voice assistants are programmed to replicate human-like interactions, a feature that often results in inflated anthropomorphization and the attribution of human qualities to these AI systems. For instance, users may opt to alter the accent or gender of their Google voice assistant's synthesized voice to create a more comforting or familiar experience, particularly when utilizing voice navigation on maps. While this capability can enhance user engagement and satisfaction, it also raises ethical concerns regarding users' expectations and perceptions. Abercrombie et al. [39] have recently shown that users tend to use gendered (in fact female) pronouns to refer to the most popular voice assistants. At the same time, technology companies routinely deny that their conversational agents are gendered or human-like, even if there is ample evidence to the contrary [40]. Cowan et al. [41] highlight the drawbacks of relying on "humanness" as a metaphor for interaction with voice assistants. The study found that many users viewed the human-like traits of voice assistants as a failed attempt to make them more relatable.

So, at least sometimes, humans interact with voice assistants employing social rules and treating them as social entities. Of course, that is not exclusive; we often treat them as objects. But in many of our interactions with voice assistants, we take the participant stance [42]. This may partly explain why the trust/trustworthiness narrative is particularly prevalent in the literature on voice assistants. In some ways, they are paradigmatic *social interaction machines* and it is possible to imagine having a meaningful relationship with them. In fact, some are proposing models of creating trust between humans and voice assistants using human-computer interaction theories and para-social relationship theory [43].

### C. Uncertainty, Assertiveness, and Transparency

When a voice assistant interacts with a human, it does so typically in the form of assertions, i.e. statements that do not indicate any degree of uncertainty. While event-driven information fusion has potential to convey degrees of confidence by voice assistants [44], users often lack context-sensitive information during interactions with these systems. Presently, there are few mechanisms in place to communicate the confidence level of a voice assistant's decision-making process to the user. For example, users may not know the probability of a specific road closure due to flash floods in a particularly vulnerable area.

Furthermore, the issue of evaluating absolute accuracy in voice assistants further complicates the matter. Hong et al. [45] conducted a comparison of four major voice assistants' performance in responding to questions about

cancer screening, revealing significant disparities and areas for improvement in the information provided. Their findings suggest issues regarding the reliability of voice assistants, particularly when sourcing responses from the internet without verifying their accuracy. This reliance on potentially inaccurate information may not only undermine what is considered to be the trustworthiness of voice assistants, but also compromise the integrity of the information they provide.

The reliance on internet-derived responses to evaluate the performance of voice assistants clearly raises pertinent concerns regarding their trustworthiness. However, this conflicts with what might be considered as the knowledge base of the voice assistant. Are we subject to trusting the voice assistant or the information available already on the web? This prompts a critical reassessment of the factors that underpin trust in AI systems, emphasizing the pivotal roles of accuracy and reliability in shaping user perceptions. Although matters like transparency remain integral to responsible AI utilization, ensuring the accuracy of information provided by voice assistants takes precedence. Users must trust the reliability of the information conveyed, necessitating the addressing of both epistemic and aleatoric uncertainties by voice assistants.

As mentioned above, transparency is ever so present in the discussion surrounding AI. Schmidt et al. [46] challenge the prevailing notion in explainable AI research that maximal algorithmic transparency inherently fosters trust. Their study underscores instances where transparency, despite its association with AI prediction accuracy, can paradoxically lead to mistrust. This finding prompts a reevaluation of the relationship between transparency and trust within the context of AI systems. Rather than viewing transparency as an absolute solution, it should be regarded as an aspect of responsible AI use. Responsible AI use entails optimizing system performance while meticulously managing confidence levels in predictions (i.e. levels of epistemic uncertainty).

Moreover, the quest for complete transparency may inadvertently create a sense of vulnerability where none should exist. Users rely on the information provided by voice assistants to make informed decisions, and any ambiguity or withholding of information regarding the uncertainty associated with suggestions undermines the reliability of the assistant. This suggests that transparency alone is insufficient to ensure the reliability of AI systems; effective communication of uncertainty is equally crucial. Users should be empowered to navigate the uncertainties inherent in AI systems, enabling them to make informed decisions based on the available information and the assistant's suggestions. Thus, the reliability of a voice assistant hinges not only on transparency but also on the clarity and completeness of the information conveyed to the user.

## IV. NEGATIVE IMPLICATIONS OF THE TRUST NARRATIVE

Now that we have presented a number of factors that contribute to the trust narrative it is important to consider

the stakes. One could reasonably ask "what's the problem?" Even if the narrative of trust and trustworthiness towards machines is misconstrued and an alternative demand for reliance would be more true to the nature of our relationship with technology, what is the harm?

Let's revisit the case of the live navigation voice assistant from the beginning of the paper. When the live navigation voice assistant led the unlucky couple into danger, did it fail at being a good navigation app? Probably not. The app is meant to provide route planning and navigation from point A to B, not guarantee a safe passage through the city. Yet, it is plausible that people implicitly rely on the navigation app for other tasks, because the interaction with it is richer than just navigation. It provides alternative routes, informs you about closing times of shops, possible road closings, natural disasters or emergencies in the area, to name a few. It is only natural, then, that people would develop a social-like relationship with their navigation voice assistant, which goes beyond treating it as an object.

#### A. Ambiguity of Attitudes towards AI

As the case demonstrates, we develop ambiguous attitudes towards AI. On the one hand, we usually treat them as objects, applying to them rules and expectations we normally apply to objects in the world. On the other hand, we sometimes extend to them social expectations and rules normally reserved for humans. These ambiguous attitudes are -at least partly- the result of the three factors we discussed in the previous section. The diverse modalities and complexity of voice assistants, in combination with the wide range of tasks for which humans use them promote a complex relationship between user and machine which goes beyond mere reliance and towards trust. At the same time, the indisputable social benefits of voice assistants make them ideal social interaction partners, moving users naturally towards the participant stance. Finally, the assertiveness and relative opacity of voice assistants restrict the freedom of decision on the part of the user. Since no degree of nuance or uncertainty is revealed, one must either accept or reject what a voice assistant says. A perception of trustworthiness is necessary for choosing to use the technology in the first place.

#### B. Bias and User Preferences

Apart from the ambiguity of attitudes towards AI, the trust/trustworthiness narrative around voice assistants may also lead to a significant conflict of principles. Because of their strong social dimension, voice assistants learn and adapt to the user's individual preferences. Over time, a user is more likely to perceive a voice assistant as trustworthy if it becomes more customised upon their beliefs, desires, and even mannerisms. However, most people hold explicit or implicit negative biases against others. So, in the process of catering to individual preferences, a voice assistant runs the risk of producing biased outputs, thus reinforcing negative stereotypes.

#### C. Mystification of AI

The final implication of the trust/trustworthiness narrative is the mystification of AI. The phenomenon is not exclusive to voice assistants, of course. However, it is important to mention that the mystification of AI often leads to unnecessary anthropomorphism, especially in the form of psychological sophistication. Perhaps the easiest way to substantiate the claim that AI systems can be trusted and trustworthy is by attributing to them agential features otherwise reserved to humans, such as agency, subjectivity, and moral status. Regardless of what one thinks about the possibility of AI agency, it is premature to assign such features to them without considering the implications. For instance, the mystification of AI could be the cause of the phenomenon of responsibility gaps [47], [48].

### V. CONCLUSIONS

#### A. Reflecting on AI Narratives: Shifting Perspectives for Future Development

While our exploration has illuminated various factors influencing the trust narrative surrounding voice-assisted AI while advocating for a shift in the narrative, it is crucial to acknowledge that we are still navigating the evolving landscape of AI technology. We recognize that surrounding concepts like trustworthy AI there is a tendency to personify AI, or a deeply rooted association to cultural pillars such as Isaac Asimov's "Three Laws of Robotics" which potentially offer thought-provoking ideals for guiding AI development. However, they remain largely in the realm of science fiction due to the current limitations of AI systems. This realization prompts us to engage in discussions that bridge the gap between theoretical aspirations and practical implementation. These discussions should be driven by the diversity in voice assistant models, the social aspect of voice assistants in human-computer interactions and the uncertainty involved in communication with voice assistants or the transparency involved in their use. As we contemplate the future of voice-assisted AI, it is essential to approach these conversations with a blend of optimism and pragmatism, recognizing both the potential and the challenges inherent in the ongoing advancement of AI technology.

### REFERENCES

- [1] New York Post - Yaron Steinbuch. *LA couple claims Google Maps sent them into dangerous South Africa 'Hell Run' area where they were attacked at gunpoint*. [Accessed 3 Feb. 2024]. [Online]. Available from: [www.nb.gy/swjeur](http://www.nb.gy/swjeur).
- [2] Baier AC. Trusting people. *Philosophical Perspectives*. 1992;6:137-53.
- [3] Becker LC. Trust as noncognitive security about motives. *Ethics*. 1996;107(1):43-61.
- [4] Holton R. Deciding to trust, coming to believe. *Australasian journal of philosophy*. 1994;72(1):63-76.
- [5] Strawson P. Freedom and Resentment. *Proceedings of the British Academy*. 1962;48:187-211.
- [6] Hawley K. Trust, distrust and commitment. *Noûs*. 2014;48(1):1-20.
- [7] Hawley K. *How to be trustworthy*. Oxford University Press, USA; 2019.
- [8] Thiebes S, Lins S, Sunyaev A. Trustworthy artificial intelligence. *Electronic Markets*. 2021;31:447-64.
- [9] Kaur D, Uslu S, Rittichier KJ, Durresti A. Trustworthy artificial intelligence: a review. *ACM Computing Surveys (CSUR)*. 2022;55(2):1-38.

- [10] Laux J, Wachter S, Mittelstadt B. Trustworthy artificial intelligence and the European Union AI act: On the conflation of trustworthiness and acceptability of risk. *Regulation & Governance*. 2024;18(1):3-32.
- [11] Ryan M. In AI we trust: ethics, artificial intelligence, and reliability. *Science and Engineering Ethics*. 2020;26(5):2749-67.
- [12] Glikson E, Woolley AW. Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*. 2020;14(2):627-60.
- [13] Choung H, David P, Ross A. Trust and ethics in AI. *AI & SOCIETY*. 2022 05;38:1-13.
- [14] Omrani N, Rivieccio G, Fiore U, Schiavone F, Agreda SG. To trust or not to trust? An assessment of trust in AI-based systems: Concerns, ethics and contexts. *Technological Forecasting and Social Change*. 2022;181:121763.
- [15] Reinhardt K. Trust and trustworthiness in AI ethics. *AI and Ethics*. 2023;3(3):735-44.
- [16] Toreini E, Aitken M, Coopamootoo K, Elliott K, Zelaya CG, Van Moorsel A. The relationship between trust in AI and trustworthy machine learning technologies. In: *Proceedings of the 2020 conference on fairness, accountability, and transparency*; 2020. p. 272-83.
- [17] Coeckelbergh M. Can we trust robots? *Ethics and information technology*. 2012;14:53-60.
- [18] Sullins JP. *Trust in robots*. The Routledge Handbook of Trust and Philosophy New York and London: Routledge. 2020.
- [19] Visser R, Peters TM, Scharlau I, Hammer B. Trust, distrust, and appropriate reliance in (X) AI: a survey of empirical evaluation of user trust. *arXiv preprint arXiv:231202034*. 2023.
- [20] Schelble BG, Lopez J, Textor C, Zhang R, McNeese NJ, Pak R, et al. Towards ethical AI: Empirically investigating dimensions of AI ethics, trust repair, and performance in human-AI teaming. *Human Factors*. 2022:00187208221116952.
- [21] Buçinca Z, Malaya MB, Gajos KZ. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*. 2021;5(CSCW1):1-21.
- [22] Wienrich C, Reitelbach C, Carolus A. The trustworthiness of voice assistants in the context of healthcare investigating the effect of perceived expertise on the trustworthiness of voice assistants, providers, data receivers, and automatic speech recognition. *Frontiers in Computer Science*. 2021;3:685250.
- [23] Kästner L, Langer M, Lazar V, Schomäcker A, Speith T, Sterz S. On the relation of trust and explainability: Why to engineer for trustworthiness. In: *2021 IEEE 29th International Requirements Engineering Conference Workshops (REW)*. IEEE; 2021. p. 169-75.
- [24] Kerasidou CX, Kerasidou A, Buscher M, Wilkinson S. Before and beyond trust: reliance in medical AI. *Journal of medical ethics*. 2021.
- [25] Benda NC, Novak LL, Reale C, Ancker JS. Trust in AI: why we should be designing for APPROPRIATE reliance. *Journal of the American Medical Informatics Association*. 2022;29(1):207-12.
- [26] López-Espejo I, Tan ZH, Hansen JH, Jensen J. Deep spoken keyword spotting: An overview. *IEEE Access*. 2021;10:4169-99.
- [27] Haeb-Umbach R, Watanabe S, Nakatani T, Bacchiani M, Hoffmeister B, Seltzer ML, et al. Speech processing for digital home assistants: Combining signal processing with deep-learning techniques. *IEEE Signal processing magazine*. 2019;36(6):111-24.
- [28] Yadav A, Vishwakarma DK. Sentiment analysis using deep learning architectures: a review. *Artificial Intelligence Review*. 2020;53(6):4335-85.
- [29] Kumar Y, Koul A, Singh C. A deep learning approaches in text-to-speech system: A systematic review and recent research perspective. *Multimedia Tools and Applications*. 2023;82(10):15171-97.
- [30] Kim S, et al. Exploring how older adults use a smart speaker-based voice assistant in their first interactions: Qualitative study. *JMIR mHealth and uHealth*. 2021;9(1):e20427.
- [31] Williams J, Azim T, Piskopani AM, Chamberlain A, Zhang S. Socio-Technical Trust For Multi-Modal Hearing Assistive Technology. In: *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*. IEEE; 2023. p. 1-5.
- [32] Elghaish F, Chauhan JK, Matarneh S, Rahimian FP, Hosseini MR. Artificial intelligence-based voice assistant for BIM data management. *Automation in Construction*. 2022;140:104320.
- [33] De Melo CM, Kim K, Norouzi N, Bruder G, Welch G. Reducing cognitive load and improving warfighter problem solving with intelligent virtual assistants. *Frontiers in psychology*. 2020;11:554706.
- [34] Guzman AL, Lewis SC. Artificial intelligence and communication: A human-machine communication research agenda. *New media & society*. 2020;22(1):70-86.
- [35] Lee K, Lee KY, Sheehan L. Hey Alexa! A magic spell of social glue?: Sharing a smart voice assistant speaker and its impact on users' perception of group harmony. *Information Systems Frontiers*. 2020;22:563-83.
- [36] McLean G, Osei-Frimpong K. Hey Alexa... examine the variables influencing the use of artificial intelligent in-home voice assistants. *Computers in Human Behavior*. 2019;99:28-37.
- [37] Guzman AL. Voices in and of the machine: Source orientation toward mobile virtual assistants. *Computers in Human Behavior*. 2019;90:343-50.
- [38] Feng H, Fawaz K, Shin KG. Continuous authentication for voice assistants. In: *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*; 2017. p. 343-55.
- [39] Abercrombie G, Curry AC, Pandya M, Rieser V. Alexa, Google, Siri: What are your pronouns? Gender and anthropomorphism in the design and perception of conversational assistants. *arXiv preprint arXiv:210602578*. 2021.
- [40] Curry AC, Rieser V. # MeToo Alexa: how conversational systems respond to sexual harassment. In: *Proceedings of the second ael workshop on ethics in natural language processing*; 2018. p. 7-14.
- [41] Cowan BR, Pantidi N, Coyle D, Morrissey K, Clarke P, Al-Shehri S, et al. "What can i help you with?" infrequent users' experiences of intelligent personal assistants. In: *Proceedings of the 19th international conference on human-computer interaction with mobile devices and services*; 2017. p. 1-12.
- [42] Marusic B. Trust, reliance and the participant stance. *Philosophers Imprint*. 2017;17(17):1-10.
- [43] Pitardi V, Marriott HR. Alexa, she's not human but... Unveiling the drivers of consumers' trust in voice-based artificial intelligence. *Psychology & Marketing*. 2021;38(4):626-42.
- [44] Zhang Y, Jiang C, Yue B, Wan J, Guizani M. Information fusion for edge intelligence: A survey. *Information Fusion*. 2022;81:171-86.
- [45] Hong G, Folcarelli A, Less J, Wang C, Erbas N, Lin S. Voice assistants and cancer screening: a comparison of Alexa, Siri, google as-

- sistant, and Cortana. *The Annals of Family Medicine*. 2021;19(5):447-9.
- [46] Schmidt P, Biessmann F, Teubner T. Transparency and trust in artificial intelligence systems. *Journal of Decision Systems*. 2020;29(4):260-78.
- [47] Matthias A. The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and information technology*. 2004;6:175-83.
- [48] Nyholm S. Attributing agency to automated systems: Reflections on human–robot collaborations and responsibility-loci. *Science and engineering ethics*. 2018;24(4):1201-19.