



King's Research Portal

DOI:

[10.1145/3543829.3544521](https://doi.org/10.1145/3543829.3544521)

Document Version

Publisher's PDF, also known as Version of record

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Seymour, W., Coté, M., & Such, J. (2022). Can you meaningfully consent in eight seconds? Identifying Ethical Issues with Verbal Consent for Voice Assistants. In *Proceedings of the 4th International Conference on Conversational User Interfaces, CUI 2022* Article 15 (ACM International Conference Proceeding Series). ACM. <https://doi.org/10.1145/3543829.3544521>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Can you meaningfully consent in eight seconds? Identifying Ethical Issues with Verbal Consent for Voice Assistants

William Seymour
william.1.seymour@kcl.ac.uk
King's College London
London, UK

Mark Coté
mark.cote@kcl.ac.uk
King's College London
London, UK

Jose Such
jose.such@kcl.ac.uk
King's College London
London, UK

ABSTRACT

Determining how voice assistants should broker consent to share data with third party software has proven to be a complex problem. Devices often require users to switch to companion smartphone apps in order to navigate permissions menus for their otherwise hands-free voice assistant. More in line with smartphone app stores, Alexa now offers “voice-forward consent”, allowing users to grant skills access to personal data mid-conversation using speech. While more usable and convenient than opening a companion app, asking for consent ‘on the fly’ can undermine several concepts core to the informed consent process. The intangible nature of voice interfaces further blurs the boundary between parts of an interaction controlled by third-party developers from the underlying platforms. This provocation paper highlights key issues with current verbal consent implementations, outlines directions for potential solutions, and presents five open questions to the research community. In so doing, we hope to help shape the development of usable and effective verbal consent for voice assistants and similar conversational user interfaces.

CCS CONCEPTS

• **Human-centered computing** → Ubiquitous and mobile computing design and evaluation methods; Natural language interfaces; • **Security and privacy** → Usability in security and privacy.

KEYWORDS

Voice assistants; consent; data protection; data sharing.

ACM Reference Format:

William Seymour, Mark Coté, and Jose Such. 2022. Can you meaningfully consent in eight seconds? Identifying Ethical Issues with Verbal Consent for Voice Assistants. In *4th Conference on Conversational User Interfaces (CUI 2022)*, July 26–28, 2022, Glasgow, United Kingdom. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3543829.3544521>

1 INTRODUCTION

A vital aspect of the software marketplaces available on today’s smart devices, including voice assistants (VAs), is the way that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CUI 2022, July 26–28, 2022, Glasgow, United Kingdom

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-9739-1/22/07...\$15.00
<https://doi.org/10.1145/3543829.3544521>

sharing of personal data with third parties is managed. Consent has emerged as the primary means of managing this relationship with skill (also called voice applications) developers and is intended to allow users to decide for themselves what they are willing to share when using third party skills. In fact, users are known to have rather different data sharing preferences depending on the third party skill [3]. The consent process is further enshrined in the GDPR, and even when skills use other legal bases for data collection, consent remains a vital *ethical* component of the design and operation of these platforms. The growing presence of VAs that are able to listen to everything said around them in phones, speakers, TVs, headphones, and other devices brings into scope wider concerns from the Ubicomp community around informed consent in the context of ubiquitous computing [10].

Unsurprisingly, designing mechanisms that broker consent for data sharing between users of voice assistants and third-party skills has proven to be a challenging problem from an HCI and UX perspective. Unlike apps which request access to device *functionality* (e.g. location determined by GPS), skills generally request access to personal data directly (e.g. the address where the device is located).¹ Currently, Alexa and Google Assistant direct users to companion smartphone apps to give consent for data sharing. This mirrors the flow of managing permissions on smartphone apps that was once commonplace, where consent would be granted in the app store at the threshold of installation. But directing users of a supposedly hands-free device to manually interact with a smartphone leads to a poor user experience and negates one of the main selling points of VAs. Last year, Amazon introduced “voice-forward consent” (VFC) for Alexa, allowing developers to request consent for data sharing mid-conversation using speech. On the face of it this not only reduces friction in the user experience, but also mirrors the shift seen in smartphone apps towards asking for permissions ‘just in time’ (a move recommended by some privacy scholars [13] but whose implementation has been criticised by others [14]).

But this current implementation of voice-based consent for VAs has several key problems and neglects core principles of informed consent. This provocation paper draws on literature from HCI and Ubicomp, as well as key data protection regulations such as the European General Data Protection Regulation (GDPR) [1] to map out concerns with current verbal consent in voice assistants, explore potential solutions, and lay out the open questions in this space for future discussion. In so doing, we hope to foster discussion within the CUI community about how VAs and similar systems might better be able to navigate consent in the future.

¹The current set of Alexa permissions are: Device Address, Customer Name, Customer Email Address, Customer Phone Number, Lists Read/Write, Amazon Pay, Reminders, Location Services, and Skill Personalisation.

User Alexa	Alexa, open Ride Hailer. Welcome to Ride Hailer. Where would you like to go?
User Alexa	The Space Needle. Sure. I need access to your name, current location, and mobile number so that I can find a ride for you.
Alexa (OS)	Do you give Ride Hailer permission to access your name, current location, and mobile number? You can say ‘I approve’ or ‘no’.
User Alexa	I approve. Thank you. A ride to the Space Needle from your current location will cost fifteen dollars, and the driver can pick you up in ten minutes. Do you want me to book it?
User Alexa	Yes. Great. Your driver will arrive in ten minutes.

Table 1: Sample VFC Flow from <https://developer.amazon.com/en-US/docs/alexa/custom-skills/use-voice-forward-consent.html>

2 CAN YOU MEANINGFULLY CONSENT IN 8 SECONDS?

Analysis of the relevant literature and regulation suggests several key issues with the voice forward consent process which we summarise here. These are: (1) the introduction of time pressure to the permissions process; (2) difficulty conveying the required amount of information via speech; (3) separating the process of granting consent from its revocation; and (4) the lack of distinction between speech from third party skills and the Alexa Operating System (OS).² Each violates established principles of informed consent, albeit in different ways.

2.1 Time Pressure

VFC’s placement within both the Alexa turn-taking model and the interaction itself adds a previously unseen sense of time pressure to the consent process. By default, Alexa times out and re-prompts the user after eight seconds, conditioning users to provide a response within this window to prevent follow-up by the device. This limits the time for decision-making, especially if the justification for data sharing by the skill developer is vague or non-existent. At the same time, the position of VFC within an interaction (as opposed to at the threshold of use) means that skills can increase the likelihood of users consenting through loss aversion [4], as refusal to grant consent will most likely result in the end of the interaction. Given this ‘take it or leave it’ Hobson’s choice after the investment of time and effort, the default response from many users is to grant consent to prevent having to repeat the interaction with a different skill, threatening the freely given nature of informed consent [15]. This is further complicated by the use of specially crafted consent dialogues on the web that are designed to steer users towards granting consent [9], with Amazon and skill designers facing a similar conflict of interest for VFC.

²The documentation refers to it as the Alexa Skill, but this terminology quickly becomes confusing.

2.2 Limited Information

The need to deliver short, concise messages in conversations causes further problems, as UX concerns preclude platforms providing information about the scope of access and associated data rights; the lower bandwidth of speech compared to graphical user interfaces means that VFC needs to accomplish in a few sentences what would normally occupy paragraphs of text. The sample VFC flow shown on the previous page highlights the brevity employed, mentioning only the name of the skill and permissions requested despite taking around twenty seconds to deliver. Information about the skill developer is omitted and it is left to the skill to accurately describe (or not) what the information will be used for, which is problematic given the known lack of traceability in skill privacy policies [6, 7]. As with the existing graphical VA consent dialogues, it is also not clear whether the skill is requesting access to data at the level of an individual person (via an associated voice profile), or the Amazon account.

2.3 Breaking Interface Symmetry

The main permissions flow, where users are directed to grant permissions in the VA’s companion app, also shows a list of all permissions that could be requested by the skill and their status (granted/not granted). This signposts to users that they can revoke consent in the future by returning to the same location in the interface. The addition of VFC moves the *granting* of consent to the voice interface, meaning that users then give consent for data sharing without seeing or being informed about how they can withdraw consent (or even that they can withdraw consent at all), as there is no mechanism for revoking consent using speech. This becomes increasingly important with the introduction of proactive modes of operation to the platform (such as ‘reminders’ and ‘routines’ that can be triggered at certain times of the day) and early work on proactive assistants [5] suggests that VAs of the future may increasingly act of their own volition. These features normalise skills running without user intervention but with the same set of permissions and ability to access personal data, further removing the experience of using a skill from the consent process.

2.4 Lack of Distinction Between System and Skills

As a final example, though delivery of voice-forward consent is handled by the Alexa OS (see Table 1) there is no audible difference between speech delivered by the Alexa OS and speech delivered by a third party skill. Users are already known to struggle when distinguishing between Alexa and third-party skills [2, 11], making it difficult for users to accurately interpret and place their trust in VFC dialogues; many will assume that Alexa controls the entire process and oversees the entire consent flow including reasons for processing data, whereas in reality skills are encouraged to account for their own usage of personal data before initiating VFC. In an extreme case, it is possible for a skill to perfectly imitate the language and response choices used in VFC without ever delegating to Alexa and with no audible difference to the user. This potential to frame consent decisions with information incorrectly attributed to the VA vendor also carries with it the risk of being exploited by skill developers to increase consent grant rates. While using

different voices to communicate speech from first and third parties is an intuitive solution, work on the *computers are social actors* paradigm suggests that people may then interpret each voice as an independent assistant [12].

3 POTENTIAL SOLUTIONS

Using the same resources used to highlight the problems with VFC we are exploring a variety of measures and mitigations that can be used by skill and platform developers in their products. At a basic level this involves strengthening policies and mechanisms that are already present, such as incorporating developer justification for using data into the skill certification process and making their disclosure a mandatory part of VFC flows.

For developers, the key first step is to motivate the need for access to data in the precious few words that are available when communicating permission requests to users. While not suggested in the Amazon developer documentation, utilising other media such as response cards and external communications channels can help to overcome the bare-bones nature of VA conversational interactions.

For VA platforms, allowing users to easily distinguish between speech from the operating system and speech from third party skills is a crucial first step that facilitates the formation of more accurate mental models by users. Modification of the usual conversational turn-taking model to facilitate engagement with consent and removal of the timeout window could help lessen feelings of being rushed. However, research on smartphone permissions suggests that in-context requests increase permission grant rates (especially where no or poor justification is given) [8], offering justification for returning to an up-front permissions decision. Pointers to additional resources, perhaps via the companion app, would give users material they can return to after the interaction is complete to foster understanding and reflection. Allowing users to revoke consent via speech, and describing how to do so when it is granted, would remove barriers to engagement by restoring the symmetry of consent flows. Where users do revoke consent, platforms could recommend alternative skills that do not require as many permissions.

Finally, the situation also presents opportunities for VAs to embrace better consent practices that will, we believe, also help people feel more confident about using devices that are often perceived as inherently unsettling. Examples of this include the potential to check-in with users in the weeks and months following an initial granting of consent, utilising the lightweight nature of VFC as a way of embracing consent as a living, ongoing practice [10]. The necessary conciseness of speech, while appearing on the surface to be a hindrance actually interacts positively with recommendations to provide “short, specific privacy notices” [13].

4 OPEN QUESTIONS

Having outlined the major issues facing voice-based consent alongside potential solutions, this section describes the open questions to the CUI community that we hope and believe that future research can address.

- (1) *How to develop the user experience of VFC?*

Usability remains a key concern in the aftermath of several

well-intentioned but ultimately poorly implemented EU privacy regulations³ and it is vital that voice-based consent does not become another burden to which people become habituated into dismissing without a second thought. This could involve user studies evaluating different implementations of the guidelines, co-design opportunities with different user groups, and policy recommendations to regulators. We envisage that a combination of these methods will ultimately be required to put voice-based consent on a firm footing for the future as VA technology develops and encompasses new aspects of daily life.

- (2) *What aspects of VFC are the most important and/or relevant?*

There are many different components to VFC, including many that are not currently implemented. The interactional and UX constraints imposed by the fundamental design of VAs means that a key step in the research agenda is determining which of these are relevant to VFC and the context in which it is used (e.g. not all aspects of informed consent, dynamic consent, or consent as a legal ground for data collection are required for VAs in the home). Prioritising how these elements are woven into the consent flow is also important, as certain aspects of VFC are limited resources (e.g. the amount of text that is delivered before a user makes a consent decision). Similarly, decisions need to be taken around which aspects are featured more prominently and which are relegated to secondary devices such as smartphones.

- (3) *How could VFC best utilise other interaction modalities?*

Related to the above, how might consent practices go beyond their current compulsory use of companion apps to best make use of connected screens, devices, and web apps to provide optional extensions to the experience? While it is true that the conversational components of VFC must be carefully balanced to remain effective and usable, the ability to back this up with text, images, augmented reality, and other experiences with devices commonly found in the home should be explored.

- (4) *How could better VFC practice be integrated into platform APIs and flows?*

Within the frameworks and architectures that currently exist, how could voice-based consent be better integrated? This could include developer justification for data collection, allowing multiple devices to be used, or more granular consent options as seen in dynamic consent. This type of work can be tricky, as the internal structures and policies of platforms is often hidden, but recent work has shown how large scale investigations into VA platform black boxes can be successful [16].

- (5) *What could change in terms of platform architecture and regulation?*

Given the continual evolution of voice assistant norms and underlying technology, as well as the gradually changing stance towards tech regulation, the next 5–10 years will undoubtedly herald changes to the market that will require re-thinking the trajectory of voice-based consent. Anticipating these changes and designing consent mechanisms

³Such as for cookies, and more recently around the GDPR.

and agendas that can be flexible in the face of uncertainty should be a high priority. This will often involve thinking and designing outside of current architectures; what communities, tools, and practices around VFC might emerge, for example, if developers had much more freedom to shape the consent experience (as opposed to having mechanisms and guidance thrust upon them) and/or the regulatory environment that was much more proactive (as opposed to cases being infrequent and mainly high profile).

5 CONCLUSION

The introduction of voice-forward consent for Alexa represents a great improvement in user experience but one that potentially undermines the informed consent process. We set out a research agenda for developing usable, effective voice-based consent mechanisms. Our own early-stage work involves using literature from HCI and Ubicomp, alongside extracts from the GDPR to highlight specific problems with the current VFC implementation and potential solutions. As this work progresses we will draw on expert opinion to further motivate and refine the ideas presented in this position paper, ultimately creating guidelines for developers and platforms that promote healthy and dynamic verbal consent practices in voice assistants and similar conversational interfaces.

ACKNOWLEDGMENTS

This work is funded by the Secure AI Assistants project via Grant EP/T026723/1 from the UK Engineering and Physical Sciences Research Council.

REFERENCES

- [1] 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *OJ L* 119 (2016), 1–8. <http://data.europa.eu/eli/reg/2016/679/oj>
- [2] Noura Abdi, Kopo M Ramokapane, and Jose Such. 2019. More than smart speakers: security and privacy perceptions of smart home personal assistants. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. 451–466.
- [3] Noura Abdi, Xiao Zhan, Kopo M Ramokapane, and Jose Such. 2021. Privacy Norms for Smart Home Personal Assistants. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [4] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. 2017. Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online. *ACM Comput. Surv.* 50, 3, Article 44 (aug 2017), 41 pages. DOI : <http://dx.doi.org/10.1145/3054926>
- [5] Narae Cha, Auk Kim, Cheul Young Park, Soowon Kang, Mingyu Park, Jae-Gil Lee, Sangsu Lee, and Uichin Lee. 2020. Hello There! Is Now a Good Time to Talk? Opportune Moments for Proactive Interactions with Smart Speakers. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 3, Article 74 (sep 2020), 28 pages. DOI : <http://dx.doi.org/10.1145/3411810>
- [6] Jide Edu, Xavier Ferrer Aran, Jose Such, and Guillermo Suarez-Tangil. 2022. Measuring Alexa Skill Privacy Practices across Three Years. In *The Web Conference (WWW)*. ACM.
- [7] Jide Edu, Xavier Ferrer-Aran, Jose Such, and Guillermo Suarez-Tangi. 2021. Skill-Vet: Automated Traceability Analysis of Amazon Alexa Skills. *IEEE Transactions on Dependable and Secure Computing (TDSC)* (2021).
- [8] Yusra Elbitar, Michael Schilling, Trung Tin Nguyen, Michael Backes, and Sven Bugiel. 2021. Explanation Beats Context: The Effect of Timing & Rationales on Users' Runtime Permission Decisions. In *30th USENIX Security Symposium (USENIX Security 21)*. 785–802.
- [9] Matthias Fassl, Lea Theresa Gröber, and Katharina Krombholz. 2021. *Stop the Consent Theater*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3411763.3451230>
- [10] Ewa Luger and Tom Rodden. 2013. An Informed View on Consent for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. Association for Computing Machinery, New York, NY, USA, 529–538. DOI : <http://dx.doi.org/10.1145/2493432.2493446>
- [11] David Major, Danny Yuxing Huang, Marshini Chetty, and Nick Feamster. 2021. Alexa, Who Am I Speaking To?: Understanding Users' Ability to Identify Third-Party Apps on Amazon Alexa. *ACM Transactions on Internet Technology (TOIT)* 22, 1 (2021), 1–22.
- [12] Clifford Nass, Jonathan Steuer, and Ellen R. Tauber. 1994. Computers Are Social Actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '94)*. Association for Computing Machinery, New York, NY, USA, 72–78. DOI : <http://dx.doi.org/10.1145/191666.191703>
- [13] Florian Schaub, Rebecca Balebako, and Lorie Faith Cranor. 2017. Designing effective privacy notices and controls. *IEEE Internet Computing* 21, 3 (2017), 70–77.
- [14] Bingyu Shen, Lili Wei, Chengcheng Xiang, Yudong Wu, Mingyao Shen, Yuan Yuan Zhou, and Xinxin Jin. 2021. Can Systems Explain Permissions Better? Understanding Users' Misperceptions under Smartphone Runtime Permission Model. In *30th USENIX Security Symposium (USENIX Security 21)*. 751–768.
- [15] Yolande Strengers, Jathan Sadowski, Zhuying Li, Anna Shimshak, and Florian 'Floyd' Mueller. 2021. What Can HCI Learn from Sexual Consent? A Feminist Process of Embodied Consent for Interactions with Emerging Technologies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 405, 13 pages. DOI : <http://dx.doi.org/10.1145/3411764.3445107>
- [16] Dawei Wang, Kai Chen, and Wei Wang. 2021. Demystifying the Vetting Process of Voice-Controlled Skills on Markets. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3, Article 130 (sep 2021), 28 pages. DOI : <http://dx.doi.org/10.1145/3478101>