Panel: Voice Assistants, UX Design and Research

Moderator:
Joseph 'Jofish' Kaye
Principal Research Scientist
Mozilla
Mountain View, CA 94022, USA
jkaye@mozilla.com

Joel Fischer
Assistant Professor
University of Nottingham
Nottingham, NG8 1BB, UK
Joel.Fischer@nottingham.ac.uk

Jason Hong
Associate Professor, CMU
5000 Forbes Ave
Pittsburgh, PA 15213-3891, USA
jasonh.cs.cmu.edu

Frank R. Bentley
Senior Principal Researcher
Yahoo
Sunnyvale, CA 94089, USA
fbentley@yahoo-inc.com

Cosmin Munteanu
Assistant Professor
University of Toronto
Ontario, CA
cosmin.munteanu@utoronto.ca

Alexis Hiniker
Assistant Professor, UW
Information School,
Box 352850, Seattle, WA 98195
alexisr@uw.edu

Janice Y. Tsai
Senior Research Scientist
Mozilla
janice@mozilla.com

Tawfiq Ammari
PhD Candidate
University of Michigan
Ann Arbor, MI, USA
tawfiqam@umich.edu

CHI'18 Extended Abstracts, April 21–26, 2018, Montréal, QC, Canada

ACM ISBN 978-1-4503-5621-3/18/04.
3/18/04.https://doi.org/10.1145/3170427.3186323

Abstract
In this panel, we discuss the challenges that are faced by HCI practitioners and researchers as they study how voice assistants (VA) are used on a daily basis. Voice has become a widespread and commercially viable interaction mechanism with the introduction of VAs such as Amazon's Alexa, Apple's Siri, the Google Assistant, and Microsoft's Cortana. Despite their prevalence, the design of VAs and their embeddedness with other personal technologies and daily routines have yet to be studied in detail. Making use of a roundtable, we will discuss these issues by providing a number of VA use scenarios that panel members will discuss. Some of the issues that researchers will discuss in this panel include: (1) obtaining VA data & privacy concerns around the processing and storage of user data; (2) the personalization of VAs and the user value derived from this interaction; and (3) the relevant UX work that reflects on the design of VAs?

Author Keywords: Voice Assistants; UX

ACM Classification Keywords
H.5.m. Miscellaneous

Introduction
New interface modalities are gaining ground as interaction mechanisms. Voice has become a

Format of Panel

Preparation: A number of use scenarios for VAs will be prepared. These will be presented by the moderator on the day of the panel

On the day of the panel: For each of the scenarios prepared, each of the panelists will have five minutes to reflect on current work or future work relevant to the scenario, specifically touching on:

- Methods used in the past or proposed in future work
- Findings from and reflections on empirical studies
- Theories that might help in grounding understanding and design of voice UX
- Privacy considerations in user personalization for voice assistants

commercially viable interaction mechanism in the last few years. Voice assistants (VAs) produced by Amazon (Alexa, Dot) and Google (Google Home) are good examples of development in this area. In this year's CES (Global Stage for Innovation), one of the biggest technology expos, voice controlled and voice activated devices produced by myriad companies made a strong showing prompting media coverage like this [1]. While this new interaction modality develops apace, the HCI community should strive to understand and enhance it.

Porcheron et al. [7] studied the situated use of VAs in multi-party conversation (in this case, they studied Siri). They found that multi-party VA use is characterized by – sometimes collaborative – repetitions and refinements of queries and mutual production of silence. In a follow-up study of Alexa in homes, the authors deployed a "conditional voice recorder" to capture conversational detail surrounding Alexa use and examine the contextual and conversational embeddedness of Alexa use [6]. In a study of Amazon reviews of Alexa, Purington et al. [8] note that the VA was not only used for accessing information or entertainment, but also as a companion for the user. They quote a reviewer writing "Alexa is my new BFF." This is by no means the only question facing the development of VAs.

Guha et al. suggest three factors for successful continued interactions with VAs: contextual assistance such as using the location of the user; updates based on user interests; and personalization, using context (defined as tasks, ongoing interests, and routines) to provide suggestions [3].

To maintain contextual assistance and continually update user interest, VAs would need to employ cloud-connected systems. Such systems introduce new challenges for maintaining users' privacy. Data, and its collection, use, and sharing are often invisible. It is very difficult to design and deploy privacy-sensitive systems employing cloud processing and storage.

Personalization in VA interactions relies heavily on machine learning (ML) techniques. Dove et al. argue for "working with ML as design material" for UX researchers [2]. This would allow UX researchers to apply their methods and understanding to the development of more usable VAs.

In our panel discussions, we will explore how researchers and industry professionals view contextualization based on user experience, privacy, VA personalization, the situated use of VA, and other issues relevant to understanding VA UX design. Following, we describe the format of this panel.

The format

The panel will begin with the moderator presenting a number of main discussion points about the use of VAs. These discussion points will be linked to use scenarios. For example, **"The user experience of older adults with VAs."** Panelists most interested in the particular scenario will each have five minutes to reflect on the scenario, design principles that might apply and its relation to their work on the topic.

The panelists will discuss how they would go about studying this particular scenario. For example:

1. What methods they would use?

2. What theoretical lenses they might rely on?
3. How we should think about UX design and research in VAs?
4. How can we best evaluate different VA designs?

Other panelists will have a chance to provide feedback should they choose to do so. We will attempt to keep the discussion time for each of the scenarios to fifteen minutes.

We will focus on three main areas of discussion: (1) VA embeddedness in daily family routines; (2) privacy concerns in in VA systems, especially with their heavy reliance on cloud services; and (3) Challenges faced by users as they interact with VAs in their daily lives. We expand on these themes below.

VA embeddedness in daily family routines. As more users interact with VA systems, the technology affects the enactment of family routines. While the embeddedness of technology has been studied in the context of technology (e.g., Rode et al. [9]), this is only starting to be the case for HCI research in the area of VAs (e.g., [6]). We will analyze the following questions among others: How do family interactions change with the introduction of more VA devices to the home? How does the introduction of VAs affect the way children interact with technology at home? How can older adults use VAs in their daily routines?

Privacy concerns in VA systems. VA systems require the use of cloud computing to analyze and personalize the use of VAs. Cloud-connected services which are "always on" present new privacy challenges to designers and researchers. Current systems based on consent cannot meet the challenges posed by VA systems [4,5]. How can we design systems that provide better privacy options for users? How can we design privacy policies more well suited for VAs and associated technological echo systems (e.g., smart lights etc.)?

VA challenges for users in their daily lives. Early research on VA suggests that users do not have a strong understanding of the way VA systems process their speech or how the queries are processed (e.g., [7]). How do users understand the way VA systems interact with them? How do users understand the capabilities and limitations of VAs? Could we provide users with a more understandable model of how these systems work?

The audience will have time to comment and ask questions of the panelists. The audience will be presented with the three main discussion areas discussed above early on in the panel discussions. The audience will have the chance to submit questions and comments to a slack channel (we might use other platforms) where other audience members can vote for specific questions/comments that will then be used to further the discussion by allowing the audience to interact with the panelists.

Organizers

Moderator: Joseph 'Jofish' Kaye is principle research scientist working in the Mozilla Emerging Technologies Team. He uses a variety of methods including big data and qualitative research to understand user needs and practices in the HCI space.

Frank Bentley is a Senior Principal Researcher at Yahoo and a Visiting Lecturer at Stanford. Frank's research explores the adoption of mobile and ubiquitous technologies using a wide variety of methods. His focus is on creating new systems based on a deep understanding of current practices and evaluating these systems in large-scale real-world deployments or products.

Prof. Cosmin Munteanu is an Assistant Professor at the Institute for Communication, Culture, Information, and Technology at University of Toronto Mississauga. His research includes speech and natural language interaction for mobile devices, mixed reality systems, learning technologies for marginalized users, usable privacy and cyber-safety, assistive technologies for older adults, and ethics in human-computer interaction research.

Joel Fischer is assistant professor (lecturer) at the University of Nottingham. He is interested in how interactive technology can be designed to support, (or disrupt) collocated people interacting face-to-face. His research approach is multidisciplinary, drawing on

ethnography, participatory design, prototyping, and studies of technology deployments.

Jason Hong is associate professor of Human-Computer Interaction Institute at Carnegie Mellon University. His work is at the intersection of usability, mobility, privacy, and security.

Alexis Hiniker is assistant professor of Human Computer Interaction (HCI) at the University of Washington School of Information. She studies how people use interactive systems, and builds new technologies to help users feel in control of their behavior and the systems they use. Her current projects investigate users' struggles to disengage from technology and the role of technology in family life.

Janice Tsai is a senior research scientist on the Mozilla Emerging Technologies team. Her research interests are in usable privacy and public policy.

Tawfiq Ammari is a PhD candidate at the University of Michigan School of Information. His research interests are at the intersection of human computer interaction (HCI) and social computing. He studies how parents, especially fathers, use social media sites to access information and social support. He is also interested in the effects of emerging technologies on family routines and interactions.

References

1. Kim Brunhuber,. 2018. The hottest thing in the world of technology: your voice. *CBC News*. Retrieved January 15, 2018 from http://www.cbc.ca/news/technology/brunhuber-ces-voice-activated-1.4483912

2. Graham Dove, Kim Halskov, Jodi Forlizzi, and John Zimmerman. 2017. UX Design Innovation: Challenges for Working with Machine Learning As a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (CHI '17), 278–288. https://doi.org/10.1145/3025453.3025739

3. Ramanathan Guha, Vineet Gupta, Vivek Raghunathan, and Ramakrishnan Srikant. 2015. User Modeling for a Personal Assistant. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining* (WSDM

’15), 275–284. https://doi.org/10.1145/2684822.2685309

4. Jason I. Hong and James A. Landay. 2004. An Architecture for Privacy-sensitive Ubiquitous Computing. In *Proceedings of the 2Nd International Conference on Mobile Systems, Applications, and Services* (MobiSys ’04), 177–189. https://doi.org/10.1145/990064.990087
5. 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), 529–538. https://doi.org/10.1145/2493432.2493446
6. Martin Porcheron, Joel E. Fischer, Stuart Reeves, and Sarah Sharples. 2018. Voice Interfaces in Everyday Life. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI ’18). https://doi.org/10.1145/2702123.2702325
7. Martin Porcheron, Joel E. Fischer, and Sarah Sharples. 2017. “Do Animals Have Accents?”: Talking with Agents in Multi-Party Conversation. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (CSCW ’17), 207–219. https://doi.org/10.1145/2998181.2998298
8. 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* (CHI EA ’17), 2853–2859. https://doi.org/10.1145/3027063.3053246
9. Jennifer A. Rode, Eleanor F. Toye, and Alan F. Blackwell. 2005. The Domestic Economy: A Broader Unit of Analysis for End User Programming. In *CHI ’05 Extended Abstracts on Human Factors in Computing Systems* (CHI EA ’05), 1757–1760. https://doi.org/10.1145/1056808.1057015