

A Mixed-Methods Approach to Understanding User Trust after Voice Assistant Failures

Amanda Baughan*
baughan@cs.washington.edu
University of Washington
Seattle, WA, USA

Xuezhi Wang
xuezhw@google.com
Google Research
New York, NY, USA

Allison Mercurio
amercurio@google.com
Google Research
Mountain View, CA, USA

Jilin Chen
jilinc@google.com
Google Research
Mountain View, CA, USA

Ariel Liu
arielliu@google.com
Google Research
Mountain View, CA, USA

Xiao Ma
xmaa@google.com
Google Research
New York, NY, USA

ABSTRACT

Despite huge gains in performance in natural language understanding via large language models in recent years, voice assistants still often fail to meet user expectations. In this study, we conducted a mixed-methods analysis of how voice assistant failures affect users' trust in their voice assistants. To illustrate how users have experienced these failures, we contribute a crowdsourced dataset of 199 voice assistant failures, categorized across 12 failure sources. Relying on interview and survey data, we find that certain failures, such as those due to overcapturing users' input, derail user trust more than others. We additionally examine how failures impact users' willingness to rely on voice assistants for future tasks. Users often stop using their voice assistants for specific tasks that result in failures for a short period of time before resuming similar usage. We demonstrate the importance of low stakes tasks, such as playing music, towards building trust after failures.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**.

KEYWORDS

voice assistants, trust, survey, interview, dataset

ACM Reference Format:

Amanda Baughan, Allison Mercurio, Ariel Liu, Xuezhi Wang, Jilin Chen, and Xiao Ma. 2023. A Mixed-Methods Approach to Understanding User Trust after Voice Assistant Failures. 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, 16 pages. <https://doi.org/10.1145/3544548.3581152>

1 INTRODUCTION

Voice assistants have received a lot of attention from both industry and academia, especially given the recent advances in natural

*This work was conducted as part of an internship with Google Research



This work is licensed under a Creative Commons Attribution International 4.0 License.

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/10.1145/3544548.3581152>

language processing (NLP). Within the past five years, advancements in NLP have achieved huge gains in accuracy when tested against standard datasets [8, 18, 34, 56, 58, 60, 63], with state-of-the-art accuracy in natural language processing models as high as 99% for certain tasks [8, 34]. This has led many practitioners and researchers alike to imagine a near future where voice assistants can be used in increasingly complex ways, including supporting healthcare tasks [41, 55], giving mental health advice [52, 66], and high stakes decision-making [17].

However, despite the increasing accuracy of NLP models and the breadth of their applications, evidences suggest that users remain reluctant and distrusting of using voice assistants [12, 29]. In the U.S., voice assistants are common in homes, with an estimated 72% of Americans having used a voice assistant [27]. However, people primarily use these for basic tasks such as playing music, setting timers, and making shopping lists [12, 29, 35]. This is because when voice assistants fail, such as by incorrectly answering a question, it derails user trust [27, 35]. User trust is pivotal to user adoption of various technologies [3], and in this case, low user trust results in reluctance to try voice assistants' novel capabilities.

As voice assistants increasingly rely on large language models [24, 47], we believe the gap between the high accuracy of these models and users' reluctance to use voice assistants for complex tasks may be explained by differences in how users and NLP practitioners evaluate the success of a model. Standard NLP models are often evaluated on large datasets of coherent text-based questions and answers [48, 49] or paired written dialogue [68]. Meanwhile, in practice users' speech may include disfluencies, such as restarts and filler words, questions not covered in training, or background noise which misconstrues speech. In the case of question answering, NLP models are evaluated based on how many questions are accurately answered on a subset of the training dataset [48, 49]. As one may expect, people can interact with voice assistants in a multitude of ways that fall outside of the scope of training data, which can lead to friction. In the eyes of users, these inaccurate responses, or voice assistant failures, can lead to frustration. For example, only five percent of users report never becoming frustrated when using voice search [13].

We believe that the gap between how NLP models are evaluated and how users encounter and perceive failures hinders the practical applications of the advancements that voice assistants have made. Therefore, we ask, which types of voice assistant failures

do users currently experience, and how do these failures affect user trust? A human-centered understanding of the types of NLP failures that occur and their impact on users trust would allow technologists to prioritize and address critical failures and enable long-term adoption of voice assistants for a wider variety of use cases.

Further, while research has started to categorize types of breakdowns in communication between users and NLP agents [28, 46], little work has looked into how users perceive these failures and subsequently trust and use their voice assistants. We draw from and extend past research to make the following contributions:

- **C1:** Iterating on the existing taxonomy of NLP failures, we crowdsource a dataset of 199 failures users have experienced across 12 different sources of failure.
- **C2:** A qualitative and quantitative evaluation on how these different failures affect user trust, specifically along dimensions of ability, benevolence, and integrity.
- **C3:** A qualitative and quantitative analysis on how trust impacts intended future use.

To accomplish this, we developed a mixed-methods, human-centered investigation into voice assistant failures. We first executed interviews with 12 voice assistant users to understand what types of failures they have experienced and how this affected their trust and subsequent use of their assistant. We concurrently crowdsourced a dataset of failures from voice assistant users on Amazon Mechanical Turk. Finally, we executed a survey to quantify how different types of failures impact users' trust in their voice assistants and their willingness to use them for various tasks in the future.

We found that different types of voice assistant failures have a differential impact on trust. Our interviews and survey revealed that participants are more forgiving of failures due to spurious triggers or ambiguity of their own request. In the case of spurious triggers, the voice assistant activates due to mishearing the activation phrase when it was not said. Users forgave this more easily, as it did not hinder them from accomplishing a goal. Failures due to ambiguity occurred when there were multiple reasonable interpretations of a request, and the response was misaligned with what the user intended while still accurately answering the question. Users tended to blame themselves for these failures. However, failures due to overcapture more severely reduced users' trust, as when the voice assistant continued listening without any additional input, users considered their use a waste of time.

We additionally find that on many occasions, users would discontinue using their voice assistant for a specific task for a short period of time following a failure, and then resume again once trust had been rebuilt. Trust was often rebuilt by using the voice assistant for tasks they considered simple, such as playing music, or alternatively, using the voice assistant for the same general task but in a different use case. In addition to these findings, we release a dataset of 199 voice assistant failures, capturing user input, voice assistant response, and the context for the failure, so that researchers may use these failures for future research on how users respond to voice assistant failures. As voice assistants continue to perform increasingly complex and high stakes tasks across various industries [17, 41, 51, 55, 66], we hope that this research will help

technologists understand, prioritize, and address natural language failures to increase and maintain user trust in voice assistants.

2 RELATED WORK

Prior research across many fields has examined the interaction between users and voice assistants, including human-computer interaction, human-centered AI, human-robotics interaction, science and technology studies (STS), computer-mediated communication (CMC), and social psychology. In addition, some work in natural language processing (NLP), especially NLP robustness, has approached technology failures in voice assistants and developed certain technical solutions to address them. Here, we provide an interdisciplinary review of research relevant to voice assistant failures during user interaction across these fields. The literature review is organized as follows: 1) literature on user expectations and trust in voice assistants; 2) human-computer interaction (HCI) approaches to understanding voice assistant failures and strategies for mitigation; 3) natural language processing (NLP) approaches to voice assistant failures, including disfluency and robustness.

2.1 User Expectations and Trust in Voice Assistants

Researchers have long tried to understand how people interact with automated agents, especially comparing and contrasting these experiences with human-to-human communication. When talking with other humans, conversations can broadly be understood as functional (also known as transactional or task-based) or social (interactional), and many conversations include a mix of both [11]. Functional conversations serve towards the pursuit of a goal, and those who participate often have understood roles towards the pursuit of that goal. In contrast, social conversations have a goal of building, strengthening, or maintaining a positive relationship with one of the participants. These social conversations can help build trust, rapport, and common ground [11].

People generally expect to have functional conversations with voice assistants [11]. The lack of social conversations may reduce users' ability to build trust in their voice assistants. Indeed, past research has shown that users trust embodied conversational agents more when they engage in small talk [5], although this varies by user personality type and level of embodiment of the agent [6]. As it stands, people report not using voice assistants for a broad range of tasks, even though they're technically capable of doing so [27]. Prior work has illustrated the importance of trust for continued voice assistant use [31, 35], as trust is pivotal to user adoption of voice assistants [33, 45] and willingness to broaden the scope of voice assistant tasks [27]. It is especially important to support trust-building between users and voice assistants as researchers continue to imagine and develop new capabilities for them, including complex tasks such as supporting healthcare tasks [41, 55], giving mental health advice [52, 66], and other high stakes decision-making [17].

This then begs the question of how trust is built between users and voice assistants. Trust in machines is an increasingly important topic, as use of automated systems is widespread [64]. Concretely, trust can be conceptualized as a combination of confidence in a system as well as willingness to act on its provided recommendations [37, 54]. Prior researchers have examined trust in machines

in terms of people’s confidence in a machine’s *ability* to perform as expected, *benevolence* (well-meaning), and *integrity* to adhere to ethical standards [36]. Broadly, past research has evaluated how various factors such as accuracy and errors affect people’s trust in algorithms [19, 20, 67]. In the case of voice assistants, Nasirian et al. [45] and Lee et al. [33] studied how quality affects trust in and adoption of voice assistants, and found that information and system quality did not impact users’ trust in a voice assistant, but interaction quality did. Interaction quality was captured based on a study by Ekinici and Dawes [22], in which Likert scale responses were captured regarding competence, attitude, service manner, and responsiveness of the voice assistant. In addition, customizing a voice assistant’s personality to the user can lead to higher trust [7], while gender does not impact users’ trust in a voice assistant [59]. Overall, prior research demonstrates the importance of the interaction quality and social conversations for building trust between users and voice assistants, which in turn affects users’ willingness to continue using them and broaden the scope of their tasks.

2.2 HCI Approaches to Voice Assistant Failures

However, there are occasionally unforeseen breaches of trust, as not all interactions go as smoothly as one expects. Prior work has explored the diversity of issues affecting engagement and ongoing use of voice assistants and has shown that when users have expectations for voice assistants that surpass its capabilities, voice assistant failures and user frustration ensues [31, 35].

This begs the question, how has prior work defined failures in voice assistants? Some work uses specific scenarios in their studies. For example, Lahoual and Frejus [31] conducted evaluation in domestic and driving situations. They identified failures due to poor voice recognition, limited understanding of a command, and connectivity. Cuadra et al. [16] used failures in specific tasks, such as attempting to give directions to an incorrect location, send a text to the wrong person, play the wrong type of music, or adding a reminder with an incorrect detail [16]. Mahmood et al. [39] simulated online shopping, in which an AI assistant with a voice component would fail by using homonyms of the requested items. For example, the ambiguous item “bow” could mean a hair bow, archery bow, or bow for gift wrapping. Salem et al. [53] had participants control a robot’s movement, and in the faulty condition, the robot would move erratically, incorrectly responding to the users’ input. Candelero et al. [9] defined failure as occasions in which someone asked a question that could not be understood or was out of scope of the voice assistants’ knowledge, in which case it would divert the conversation to ask an unrelated question.

Other research aims to provide a broad categorization of voice assistant failures, drawing from theoretical frameworks of communication between humans [10, 28, 46]. We reference Herbert Clark’s grounding model for human communication, which relies on four different levels to achieve mutual understanding: channel, signal, intention, and conversation [10]. This was expanded by Paek and Horvitz [46], which applied these four levels to human-machine interactions and failure points. Channel level errors include when an AI fails to attend to a users’ attempt to initiate communication; signal level errors include an error in capturing user input (e.g. due to transcription); intention level errors include mistakes in

making sense of the semantic meaning of the transcribed input; and conversation level errors occur when a user has requested an unknown action to the AI (e.g. asking a weather app to schedule something). Hong et al. [28] built on this model, specifically restricting the context to NLP failures, rather than AI as a whole. Based on interviews with NLP practitioners, they renamed the categories as attention (channel), perception (signal), understanding (intention), and response (conversation). Hong et al. [28] focused on failures that are either very common, or rare but very costly, to cover the most important and frequent failures users encounter when interacting with NLP-based systems. In this work, we build on their existing taxonomy of NLP failures [28], narrowing the use case to only voice assistant failures, and evaluating how different failures impact on user trust and future intended use.

There is currently little systematic evaluation of the impact of voice assistant failures on user trust. Salem et al. [53] found that if a robot had faulty performance, this did not influence participants’ decisions to comply with its requests, but it did significantly affect their perceptions of the robot’s reliability and trustworthiness. Mahmood et al. [39] found that voice assistants that accepted blame and apologized for mistakes were thought to be more intelligent, likeable, and effective in recovering from failures than assistants that shifted the blame.

Sometimes after a failure, users will try to reformulate, simplify, or hyper-enunciate their commands as a way to continue using the device [31, 35, 43, 61]. If users are repeatedly unable to repair failures with voice assistant, this weakens their trust and causes them to reduce their scope of commands to simple tasks with low risk of failure [31, 35]. Lahoual and Frejus [31] found that in some situations, voice assistant failures can erode trust to the extent that users abandon voice assistants all together. However, not all failures require self-repair. A study by Cuadra et al. [16] found that when voice assistants make mistakes, voice assistant self-repair greatly improves people’s assessment of an intelligent voice assistant, but it can have the opposite impact if no correction is needed. Thus, understanding which types of failures undermine trust the most may also inform us when failure mitigation strategies should be activated.

2.3 NLP Approaches to Voice Assistant Failures

The NLP community has also examined voice assistant failures from a slightly different angle, focusing on the robustness of different NLP components underlying voice assistants, such as models for tasks in natural language inference [44], question answering [26, 42], and speech recognition [32]. NLP robustness can be defined as understanding how model performance changes when testing on a new dataset, which has a different distribution from the dataset the model is trained on [62]. In practice, users’ real world interactions with voice assistants could differ from data used in development, which mimics the data distribution shift in NLP robustness research.

Such data distribution shifts are shown to lead to model failures. In the case of question answering, state-of-art models perform nearly at human-level for reading comprehension on standard benchmarks collected from Wikipedia [49]. However, Miller et al. [42] found that model performance drops when the question answering model is evaluated on different topic domains, such as

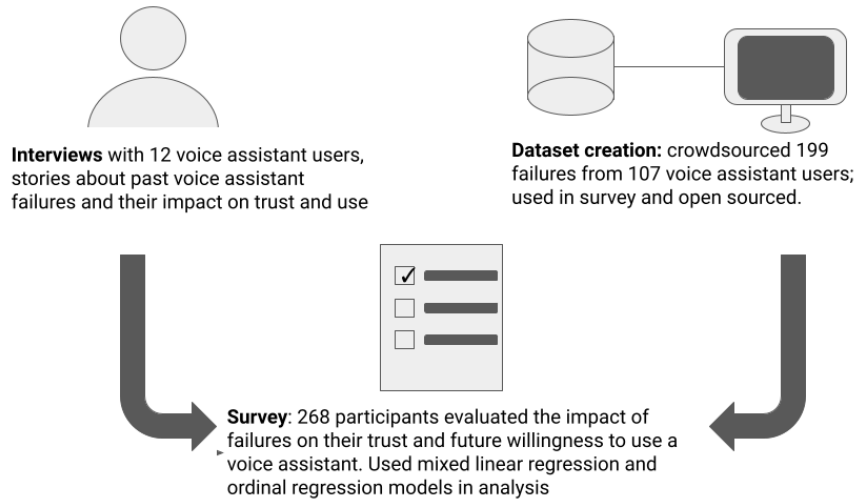


Figure 1: To analyze the impact of voice assistant failures on user trust, we used a mixed-methods approach, including interviews and a survey. As part of the materials for our survey, we crowdsourced 199 failures from 107 voice assistant users, and include this dataset as part of our contributions.

New York Times articles, Reddit posts, and Amazon product reviews. Noisy input can also harm model performances. Lee et al. [32] showed speech recognition errors have catastrophic impact on machine comprehension. Gupta et al. [26] created a question answering dataset Disflu-QA where humans introduce contextual disfluencies, which also lead to model performance drops.

Although these works do not directly focus on voice assistant failures, topic domain changes, speech recognition errors and disfluencies are all very common during user interactions with voice assistants. Such similarities motivate us to draw parallels between the NLP robustness literature and HCI perspectives of system failures. By understanding how different types of failures affect trust in voice assistants overall, we can then try to pinpoint the underlying NLP components that are the root cause of the most critical failures that erode trust [30]. Technical solutions can then be leveraged to improve the robustness of the most critical parts of the system in order to increase user trust and long-term engagement most efficiently.

3 METHOD OVERVIEW

Now that we have established the importance of understanding of how voice assistant failures impact user trust, we proceed to conduct a mixed-method study. First, to prepare for the quantitative evaluation, we reviewed existing datasets in HCI and NLP to find failures that we could use as materials for our survey. Ultimately, the existing datasets were not sufficient for our needs. Therefore, we crowdsourced a dataset of failures from voice assistant users, which we also open source as part of the contributions of this study. Concurrently, we conducted interviews with 12 voice assistant users to understand which types of failures they have experienced, and how this affected their trust in and subsequent use of the assistant.

These interviews were designed to provide a broad understanding of the thoughts, feelings, and behaviors that users have with regard to voice assistant failures and inform the quantitative survey design. Finally, we executed a survey to quantify how different types of failures impact user perceptions of trust in their voice assistants and their willingness to use them for various tasks in the future. To report these findings, we first describe our process of collecting the crowdsourced dataset of failures, and how we selected a subset to use in our survey. Next, we present the interviews and survey, first describing our data collection and analysis, and then presenting the results concurrently.

4 CROWDSOURCING A DATASET OF VOICE ASSISTANT FAILURES

The first goal in our investigation was to determine which types of failures users experience when using voice assistants. We first evaluated existing datasets for fit and breadth of failures. We determined they were not sufficient for our purposes, so we proceeded to crowdsource a dataset of failures, adapting a taxonomy from Hong et al. [28] to guide our collection. Finally, we cleaned and open-sourced this dataset as a contribution of our work.

4.1 A Review of Existing HCI and NLP Datasets

We first explored benchmark datasets in NLP, which contain a large number of either questions and answers [48–50], or conversational dialogue [25, 57, 68]. We found that existing NLP datasets do not cover the wide breadth of possible conversational failure cases due to their emphasis on correct data for training. Additionally, their focus on specific task performance, such as answering questions or dialogue generation, is more narrow than the variety of use cases for voice assistants. As training data relies on accurate task completion,

Failure Type	Sequential Coding Guide	Failure Source	Failure Scenario
Attention	A lack of visual or audio evidence the voice assistant has started listening OR video or audio evidence that the voice assistant has started listening in the absence of a cue	Missed Trigger	Users say something to trigger the voice assistant, but it fails to respond.
		Spurious Trigger	Users do not say something to trigger the voice assistant, but it activates anyways.
		Delayed Trigger	Similar to system latency, the users say something to trigger the voice assistant, but it replies too late to be useful.
Perception	Visual or audio evidence that speech is being incorrectly captured by the system. For example, being cut off by the voice assistant, witnessing it continue to listen once the users' speech is complete, clearly mishearing a word, or evidence of background noise and cross-talk.	Noisy Channel	User input is incorrectly captured due to background noise.
		Overcapture	The voice assistant captures more input than intended by either beginning to capture input too early or ending too late, and acting on external data not relevant to the users' request.
		Truncation	System does not fully capture users' speech, by either beginning to capture input too late or ending too early.
		Transcription	System generates a transcription error, often in the form of similar sounding words.
Understanding	Suspecting that audio was correctly captured but not mapped to the correct action. For example, receiving a response indicating inability to complete an action that has worked in the past, or receiving a response that is plausible but not correct for the intention of the input.	Ambiguity	There may be several interpretations of the users' intent, and the system responds in a way that is plausibly accurate but not correct for the users' intent.
		Misunderstanding	The system maps the users' input to an incorrect action, perhaps with some correct inference on the users' intent, but not fully accurate.
		No Understanding	The system fails to map the user's input to any known action or response.
Response	Finally, assuming that the input was correctly captured and understood, was the response generated incorrect, unclear, not given, or otherwise wrong?	Action Execution: No Action Action Execution: Incorrect Action	If the system listens to the full request, but then turns off before giving any type of answer or taking action. The system gives information that is incorrect.

Table 1: Qualitative codebook and description of the various failures that were collected. We checked each failure for failure type sequentially, starting by checking if it could be an attention failure and progressing through the types until we found one that fit. From failure type, we then assessed which failure source applied.

these datasets did not contain failures. While testing these models produces a small percentage of errors (roughly 10%), the types of failures could only fall in the response and understanding categories, as attention and perception failures are excluded from the context of training these types of models. This limited their usefulness for our purpose of understanding voice assistant failures that occur in use and their impact on user trust.

In addition to these benchmark datasets, we investigated datasets that incorporated spoken word speech patterns, such as the Spoken SQuAD dataset [32] and Disflu-QA dataset [26], as well as human-agent interaction datasets, such as the ACE dataset [1], the Niki and Julie corpus [2], and a video dataset of voice assistant failures [15]. In these cases, we found that the datasets were still restricted to only failures at the understanding and response level [26, 32] or the context for the failures was very specific and did not necessarily capture the breadth of possible failures users experience [1, 2]. Cuadra et al. [15]'s video dataset was the closest available fit for our needs, but we still found the use case of in-lab question-answering too narrow for our purposes. Therefore, we decided to crowdsource a dataset of voice assistant failures from users, and use these failures when conducting our quantitative survey on user trust.

4.2 Dataset Collection

4.2.1 Procedure. Crowd workers were asked to submit three failures they had experienced with a voice assistant. They were asked about three specific types of failures out of a taxonomy of 12, which were randomly chosen and displayed in equal measure across all workers. The taxonomy of failures that we used to ask about specific types of failures was adapted from previous work by Hong et al. [28], and identifies failures due to attention, perception, understanding, and response, as shown in Table 1. Each question began by asking users if they could recall a time when their voice assistant had failed, based on the definitions in our taxonomy. For example, to capture missed trigger failures we asked “Has there ever been a time when you intended to activate a voice assistant, but it did not respond?” If so, we asked these workers to include 1. what they had said to the voice assistant, 2. how the voice assistant responded, 3. the context for the failure, including what happened in the environment, and 4. the frequency at which the failure occurred from 1 (rarely when I use it) to 5 (every time I use it). These were all presented as text entry boxes except for the frequency question, which was multiple choice. Crowd workers were additionally asked to optionally share an additional failure that they had not had the

chance to share already. This was included to capture failures that did not fit any of the three the categories they were presented with, and we then categorized these failures according to our taxonomy.

Once we received these failures, we anonymized the type of voice assistant in the submitted examples, replacing activation words with “Voice Assistant” for consistency. We then edited grammatical and spelling errors for clarity. We also removed failures if they were not on-task, unclear, or exact repeats of other submitted failures. Finally, we noticed that some of the categories the users submitted the failures under were incorrect, so we re-categorized the failures according to the codebook we developed as outlined in Table 1. Two raters iteratively coded 101 submitted failures, with a final coding session achieving an interrater agreement of 70%. One researcher then went back and coded the entire dataset in its entirety. In total, our finalized dataset contains 199 failures across 12 categories, submitted by 107 unique crowd workers.

4.2.2 Crowd Worker Characteristics. We used Amazon Mechanical Turk to recruit the crowd workers. In total, 107 crowd workers contributed to our dataset. We required workers to have the following qualifications: a HIT Approval Rate over 98%, over 1000 HITs approved, AMT Masters, from the United States, over the age of 18, and voice assistant users on at least a weekly basis. The plurality of users were in the age range of 35-44 ($n = 46$), followed by 25-34 ($n = 32$), and 45-54 ($n = 16$), with the rest falling in 55-64 ($n = 8$), 18-24 ($n = 1$), and 1 preferring not to answer. Fifty-eight crowd workers were men, 44 were women, 1 preferred not to answer, and 1 identified as both a man and a woman. They used commercial voice assistants such as Amazon Alexa ($n = 59$), Google Assistant ($n = 62$), and Apple’s Siri ($n = 40$), with many using some combination of the three ($n = 47$). 91 crowd workers were native English speakers, and 13 were not. The plurality identified as White ($n = 58$), and 39 identified as Asian. Three crowd workers did not provide any demographic information. The task took 15-20 minutes to complete on average, and they received \$5.00 USD compensation.

4.2.3 Final Dataset. In total, our finalized dataset contained 199 failures from 107 users across 12 different types of failures according to the taxonomy based on Hong et al. [28], as updated in Table 1. The failures we received most often were due to misunderstanding ($n = 38$), missed trigger ($n = 25$), and noisy channel ($n = 22$). Users least often submitted failures for truncation ($n = 7$), overcapture ($n = 7$), and delayed triggers ($n = 8$). Most crowd workers submitted failures saying that they happened “rarely when I use it” ($n = 87$) or “sometimes when I use it” ($n = 84$). Example failures across the 12 categories can be found in Table 2.

On average, the highest frequency of failures occurred for no understanding ($m = 2.15$, sometimes when I use it, $sd = 0.67$) and action execution: incorrect ($m = 2.00$, sometimes when I use it, $sd = 0.88$). The rest of the failure sources had an average reported frequency between 1.0 (rarely when I use it) and 2.0 (sometimes when I use it). The lowest frequency failures were due to delayed triggers ($m = 1.25$, $sd = 0.46$) and ambiguity ($m = 1.39$, $sd = 0.78$).

We then used 60 of the failures from our dataset in our survey to quantify the impact of different failures on user trust. This is outlined in more detail in the following section. This dataset has

been open sourced¹ for researchers to use to answer future research questions related to voice assistant failures in the future.

5 INTERVIEW AND SURVEY METHODS

Once we had gathered and categorized our dataset of voice assistant failures, we were ready to answer our research question: how do voice assistant failures impact user trust? To do so, we first conducted exploratory interviews with 12 people to gather their thoughts, feelings, and behaviors after experiencing voice assistant failures. We used these findings and the failures collected in the dataset to then design and execute a survey. This quantified how various voice assistant failures impact users’ trust, as measured by their perceptions of the voice assistant’s ability, benevolence, integrity, and their willingness to use it for future tasks. Here, we describe the methods for both the interviews and survey, and we follow this by jointly presenting the results from both studies.

5.1 Interview Methods

5.1.1 Interview Procedure. Interviews began with questions about why the participants chose to start using voice assistants and what types of questions they frequently would ask of them. We asked for common times and places they would use their voice assistants to understand their general experience with voice assistants.

Once these were established, we asked participants to tell us about a time they were using their voice assistant and it made a mistake, in as much detail as they could recall. We asked what they had been trying to do and why, if others were present, and if anything else was happening in their environment. We probed for users’ feelings once the failure occurred, and their perceptions about the voice assistant’s ability to understand them and give them accurate information. We asked participants what they did in the moment to respond to the failure. Finally, we asked questions about their use of the voice assistant in the aftermath, including how much they trusted it and if they changed any of their behaviors to mitigate future failures. All interviews were conducted remotely.

5.1.2 Interview Participants. During recruitment, we asked participants to submit their demographic information, how frequently they used voice assistants and on what types of devices. We additionally required participants to write a short (1-3 sentence) summary of a time they encountered a failure while using their voice assistant. We selected participants based on demographic distribution and the level of detail they included regarding the failure.

All of our 12 participants lived in the United States. They used voice assistants at least 1-3 times a week ($n = 2$), with the majority reporting using a voice assistant every day ($n = 8$), and the rest ($n = 2$) using it 4-6 times a week. The majority of participants used a voice assistant on their mobile device ($n = 11$), and five of these participants also used a voice assistant smart home device. One participant only used a voice assistant smart home device. Participants reported using common commercial voice assistants such as Amazon Alexa ($n = 2$), Google Assistant ($n = 7$), and Apple’s Siri ($n = 8$). Participants’ ages ranged from 18 to 50, with the plurality ($n = 5$) in the age range of 18-23. 3 of our participants were 41-50, 2 were 31-40, and 2 were 24-30. Six of our participants

¹<https://www.kaggle.com/datasets/googleai/voice-assistant-failures>

Failure Source	Context	What the User Said	How the Voice Assistant Reacted
Missed Trigger	I tell her to set a timer for ten minutes, I was alone and no one present at the moment.	Voice Assistant, set a timer for 10 minutes.	[No response.]
Spurious Trigger	I was having a conference call with my team, and I was calling my coworker Sherry. The voice assistant mistakenly got turned on.	[While talking to the coworker] "Can you share your screen?"	[Responding to the conversation with the coworker.] "One moment, let me help you with that"
Delayed Trigger	It happened while I was driving a car.	Voice Assistant, show me the route to the national park.	[The voice assistant takes so much time to respond that before it can respond, you once again ask the route.]
Noisy Channel	My children were playing in the background and the dog was barking, and I had to raise my voice and try several times to be heard by my phone even though it was inches from my face.	Voice Assistant, what's the weather?	[It didn't realize that my request had ended] and kept spinning.] "I'm sorry, I didn't quite understand you."
Overcapture	I was telling it to turn off the lights. I was the only one there.	Voice Assistant, turn off the lights.	[It continues listening for so long that you turn them off yourself.]
Truncation	I asked the voice assistant to calculate a math question, but it cut me off.	Voice Assistant, can you multiply 54, 39, 33, and 22?	"54 times 39 times 33 is 69,498."
Transcription	I asked for the weather conditions in the city I live in. No others were present except for me.	Voice Assistant, what is the temperature in Murrieta, CA today?	"The temperature in Marietta, Georgia today is 65 degrees Fahrenheit."
Ambiguity	I was at home, alone, watching UFC and asked how old a fighter was.	Voice Assistant, how old is Johnny Walker?	"Johnny Walker was founded in 1865." [It referred to the whiskey company instead of the fighter.]
Misunderstanding	I asked it to play the theme from Halloween. I was sitting with my mother.	Voice Assistant, play the theme song to the movie Halloween.	[It plays a scary sounds soundtrack instead of the song.]
No Understanding	I was trying to run a routine to wake up my kids.	Voice Assistant, wake up the twins.	"Sorry, I don't know that." [However, I've set up a routine for "wake up the twins" that has worked in the past.]
Action Execution: No Action	I asked when a movie was coming out in theatres, and it kept spinning its light over and over.	Voice Assistant, when does Shang-Chi come out in theatres?	[Pauses for a really long time, then turns its lights off and does not respond.]
Action Execution: Incorrect Action	I was at home, in my living room, alone. I was trying to find out how long Taco Bell was open.	Voice Assistant, when does the Taco Bell on Glenwood close.	"Taco Bell is open until 1am". [Upon driving to Taco Bell, I realized it closed at 11:30pm.]

Table 2: Table of voice assistant failures users submitted, including the context for the failure, what the user said, and what the voice assistant said.

identified as women, five participants identified as men, and one participant identified as non-binary. Three participants identified as Asian, three identified as White, three identified as Black or African American, two identified as Hispanic, Latino, or Spanish origin, and one identified as both White and Black or African American. All of our participants spoke English as a native language. Participants were compensated with a \$50 gift card and each interview lasted roughly 30 minutes.

5.1.3 Interview Analysis. Interviews were transcribed in their entirety by an automated transcription service and analyzed via a deductive and inductive process [14]. We used deductive analysis to assess which types of failures these participants experienced. To ground our deductive analysis, we used the same codebook as we did for the dataset, as demonstrated in Table 1. We first identified instances in which participants were discussing distinct failures, and then applied our codebook to these instances. We used cues such as what was happening in their environment, and when appropriate, users' own perceptions of why the failure occurred. We began by first identifying if failures belonged in which of the four failure types: attention, perception, understanding, or response. First, to determine if there was an attention failure, we investigated if there was evidence that the voice assistant accurately responded to an activation phrase, as indicated by visual or auditory cues, or otherwise

by the participant's narrative. Second, we evaluated if there was an error in perception, based on the participants' assumption of if the voice assistant accurately parsed the input from the participant, our own assessment from their narrative, or other audio/visual cues. Next, assuming that the input was correctly parsed, we sought to understand if the voice assistant accurately understood the semantic meaning of the input (understanding failures), using the same process. Finally, assuming all else had been correctly understood, we assigned response failures, indicating that the voice assistant either did not take action or took the incorrect action in response to an accurately understood command. Once a failure type was determined, we then further specified the failure sources as noted in Table 1. We resolved disagreements both asynchronously and in meetings, through discussion and comparison, over the course of several weeks.

While conducting this analysis, we also inductively identified themes related to these failures' impact on future tasks and recovery strategies. To conduct this analysis, two researchers reviewed the twelve transcripts in their entirety, and one additional researcher reviewed five of these transcripts to further broaden and diversify themes. These researchers met over the course of several weeks to compare notes and themes, ultimately creating four different themes through inductive analysis. Of these themes, we report two

due to their novelty, specifically as related to future task orientation and recovery strategies.

5.2 Survey Methods

To quantify our findings from interviews, we developed a survey to explore users' trust in voice assistants following each of the twelve different types of failures from our taxonomy, as well as their willingness to use voice assistants for a variety of tasks in the aftermath.

5.2.1 Procedure. The survey contained a screener, the core task, and a demographic section. We required participants be over 18 years old, use their voice assistant in English, and use a voice assistant with some regularity to participate. If participants passed the screener, they were required to review and agree to a digital consent form to continue.

The core task stated, "*The following questions will ask you what you think about the abilities of a voice assistant, given that the voice assistant has made a mistake. Imagine these mistakes have been made by a voice assistant you have used before. Please consider each scenario as independent of any that come before or follow it. This survey will take approximately 20 minutes.*" Participants were then presented with 12 different failure scenarios, and they were asked to rate their trust in two separate questions.

The first question measured trust in voice assistants as a confidence score across three dimensions: ability, benevolence, and integrity. These were selected because prior work on trust has determined these elements explain a large portion of trustworthiness [36, 40]. In the context of voice assistants, ability refers to how capable the voice assistant is of accurately responding to users' input. Benevolence refers to how well-meaning the product is. And finally, integrity represents that it will adhere to ethical standards.

We asked participants to rate their confidence in voice assistants' ability, benevolence, and integrity, as a percentage on a scale of 0-100, with steps of 10, to replicate how prior work has conceptualized trust [36]. This was captured in response to the following statements:

- (Ability) This voice assistant is generally capable of accurately responding to commands.
- (Benevolence) This voice assistant is designed to satisfy the commands its users give.
- (Integrity) This voice assistant will not cause harm to its users.

The second question evaluated users' trust in the voice assistant to complete tasks that required high, medium, and low trust. To select these tasks, we ran a small survey on Mechanical Turk with 88 voice assistant users. We presented 12 different questions, which first gave an example voice assistant failure (one for each failure source), and then asked "*How much would you trust this voice assistant to do the following tasks:*" give a weather forecast, play music, edit a shopping list, text a coworker, and send money. Users could choose that they would trust it completely, trust it somewhat, or not trust it at all.

There was not a significant difference in how much people trusted the voice assistant to play music compared to forecast the weather ($Z = 2.06, p = 0.078$). There was also not a significant difference in how much people trusted the voice assistant to edit a shopping cart or text a coworker ($Z = 1.39, p = 0.21$) as determined

by pairwise comparisons, using Z -tests, corrected with Holm's sequential Bonferroni procedure on an ANOVA of an ordinal mixed model. We found that there were significant differences between playing music, texting a coworker, and transferring money, with users having the most trust in the voice assistant playing music after a failure, less trust in texting a coworker, and still less in transferring money. Therefore, we selected playing music, texting a coworker, and transferring money to represent low, medium, and high levels of trust required. Therefore, after asking about ability, benevolence, and integrity, we asked participants how much they trusted their voice assistants to execute the following tasks: play music, text a coworker, and transfer money. These questions were displayed on a linear scale of 1 ("I do not trust it at all") to 5 ("I completely trust it"), with steps of 1.

We completed the survey with an open-ended, optional question for participants to share anything else they would like to add. The survey concluded with demographic questions regarding gender, race, ethnicity, whether they were native English speakers, what type of voice assistants they used, and their general trust tendency as control variables. General trust tendency was measured based on responses to the following: "*Generally speaking, would you say that most people can be trusted, or that you need to be very careful in dealing with people?*" The options ranged from 1 (need to be very careful in dealing with people) to 5 (most people can be trusted). The questionnaire used for the survey has been submitted as supplementary materials.

5.2.2 Materials from our Dataset. To present each of the twelve failure sources in our survey, we drew from the dataset we had created. We selected five failures from each of the twelve categories. We required that these failures had been coded by two of the team members who were in agreement (see dataset examples in Table 2). We used random selection to determine which of the five possible failures was presented to each user for each failure source. These are denoted in the dataset "Survey" column.

5.2.3 Participants. We recruited participants from Amazon Mechanical Turk. We first ran a small pilot ($n = 27$) in which we determined that participants completed the survey in roughly 20 minutes on average, and we set the compensation rate at \$9 USD. After removing participants who did not pass the attention check or straight-lined, meaning they responded to every question with the same answer, we had a total of 268 participants. These participants were required to have the following qualifications: AMT Masters, with over 1000 HITs already approved, over 18 years old, live in the United States, an approval rate greater than 97%, and they must not have participated in any of our prior studies.

The plurality of our participants were in the age range of 35-44 ($n = 106$), followed by 25-34 ($n = 68$), 45-54 ($n = 52$), 55-64 ($n = 33$), with 2-4 participants in each of the age brackets of 18-24, 65-74, and 75+. 134 of our participants identified as men, 132 identified as women, and 2 identified as non-binary genders. The majority of our participants were White ($n = 210$), 21 participants were Black, and 15 were Asian. The rest of our participants identified as mixed race or preferred not to answer.

	Ability				Benevolence				Integrity			
	<i>F</i>	<i>df</i>	residuals	<i>p</i>	<i>F</i>	<i>df</i>	residuals	<i>p</i>	<i>F</i>	<i>df</i>	residuals	<i>p</i>
(Intercept)	9844.95	1	422.35	<0.001	8436.32	1	355.86	<0.001	10153.41	1	325.54	<0.001
General Trust	3.07	4	286.67	0.017	1.78	4	299.1	0.133	4.69	4	311.17	0.001
Failure Type	17.17	3	2656.78	<.001	8.87	3	2711.23	<.001	20.56	3	2772.09	<.001

Table 3: Voice assistant failure types significantly impacted users trust in voice assistants, across ability, benevolence, and integrity when controlling for their baseline trust tendencies, based on an ANOVA of three linear mixed models. Failure type was encoded as a categorical variable, and general trust was encoded as an ordinal value. Participant ID was a random, categorical variable.

6 RESULTS: TRUST IN VOICE ASSISTANTS AFTER FAILURES

In interviews, we found that participants reported failures across all four failure types and ten of the twelve failure sources. The only two failure sources that were not mentioned in interviews were missed triggers and delayed triggers in the attention failure type. To understand which types of failures most significantly impacted user trust, we analyzed how various failures impacted users’ confidence in their voice assistant’s ability, benevolence, and integrity. We used six mixed-linear regression models with log-normalized confidence in either ability, benevolence, or integrity as the numeric dependent variable. Note that there are two different levels at which we conduct the analysis. The first is at the four broad “failure types” level (attention, perception, understanding, and response). Then we drill down to the detailed 12 “failure sources” nested within each failure type. Therefore, for each dimension of trust, we encoded failure type or failure source, as well as general trust tendency, as independent variables, so there were two regression models per dimension of trust. Failure type and failure source were encoded as categorical variables, and general trust tendency was encoded as an ordinal value. In all models, PID was encoded as a random, categorical variable.

An ANOVA on the regression models revealed that failure type (attention, perception, understanding, response) does significantly impact perceptions of ability ($F(3, 2656.78) = 17.17, p < .001$), benevolence ($F(3, 2711.23) = 8.87, p < 0.001$), and integrity ($F(3, 2772.09) = 20.56, p < .001$) when controlling for general trust tendency (see Fig. 2 and Table 3). We found that the failure type “Response” (which includes action execution: inaction and action execution: incorrect action) more significantly deteriorated user trust in voice assistants across ability ($m = 43.6, \beta = -0.155, p < .001$), benevolence ($m = 52.5, \beta = -0.072, p = 0.013$), and integrity ($m = 57.3, \beta = -0.124, p < .001$), compared with failures due to “Attention” (which includes missed triggers, spurious triggers, and delayed triggers). Attention failures had a mean trust in ability of 49.9, benevolence of 56.0, and integrity of 67.7 on the scale of 0-100%. We also found that failures due to perception significantly reduced users’ confidence in voice assistant’s ability ($m = 44.8, \beta = -0.122, p < .001$) and benevolence ($m = 53.6, \beta = -0.047, p = .05$), but had no measurable effect on integrity ($m = 61.4, \beta = -0.014, p = 0.484$) compared with attention failures. Failures due to understanding maintained higher user confidence in benevolence ($m = 57.9, \beta = 0.054, p = 0.031$) and integrity ($m = 65.0, \beta = 0.063, p = 0.003$), but had no measurable

effect on ability ($m = 50.2, \beta = 0.015, p = 0.592$) compared with attention failures.

Overall, response failures had the lowest average scores across ability, perception, and integrity. The most drastic difference between these categories is between failures due to understanding, which generally maintained the highest levels of trust in ability, benevolence, and integrity, as shown in Fig. 2. Therefore, in the analysis to follow that evaluates changes in trust across failure source, *response: incorrect action* has been chosen as the reference variable, and all betas reported are in reference to this category. Below, we explore in more detail how users across both interviews and the survey responded to failures across the various failure sources.

6.1 Attention Failures

Attention failures are any failures in which a voice assistant does not accurately respond to an attempt for activation. These were the least commonly reported failures across interviews. In the survey, failures due to missed triggers were particularly harmful to users’




Figure 2: Average scores across the three dimensions of trust (ability, benevolence, and integrity) by failure type. Participants expressed higher confidence across three trust dimensions after encountering attention and understanding failures, compared to perception and response. Error bars display the confidence interval.




Figure 3: Average scores across the three dimensions of trust (ability, benevolence, and integrity) by failure type. Participants expressed higher confidence across three trust dimensions after encountering failures due to ambiguity and spurious triggers than they are of other failure types, especially missed triggers and overcapture failures.

confidence in voice assistants’ ability ($m = 41.2, \beta = -0.10, p = 0.05$) and benevolence ($m = 47.1, \beta = -0.228, p < 0.001$). However, the impact on integrity was positive compared to the reference value (action execution: incorrect action) ($m = 61.8, \beta = 0.194, p < 0.001$). None of the interview participants reported failures due to missed triggers. As all of the failures had a favorable impact on integrity compared to the reference value, we refrain from reporting it throughout the rest of the results. See Fig. 3 and Table 4 for more details.

Only P7 and P12 reported experiencing attention failures in interviews, and they were both spurious triggers. As shown in Table 1, these are failures in which the voice assistant activates in the absence of an activation phrase. P7 reported that,

“I feel like in conversation if I have it plugged in and there’s like multiple people in the room, and they’re talking or whatever, I think sometimes it may hear an [activation phrase] where it’s not. And if that’s happened, where it’s activated like once or twice completely out of nowhere, and that hasn’t upset me or anything, but it’s, it was just like, I didn’t say [an activation phrase]. Why are you activating? What’s happening? Why are you doing this?”

P7 additionally said they were working and “It must have heard an [activation phrase] somewhere in there. And then it started speaking while I was trying to do my [job], and I had to like stop and be like, hey, stop.” They said, “It would really piss [me off].” Similarly, P12

reported that these types of failures were “irritating but funny at the same time.” They said they were funny “because sometimes like, when you’re usually calling [the voice assistant] she’ll take a longer time to respond, but when you’re not talking to it, it automatically pops up... Like, I’m not talking to you, but you could answer me when I’m talking to you.” As demonstrated in Fig. 3 and Table 4, failures due to spurious triggers had a more favorable relative impact on users’ impressions of trust in the voice assistant’s ability ($m = 57.9, \beta = 0.36, p < 0.001$) and benevolence ($m = 64.2, \beta = 0.236, p < 0.001$). Overall, it appears that these are one of the least detrimental failures to users’ trust.

Similarly, failures due to delayed triggers were favorable to users’ perceptions of ability ($m = 49.3, \beta = 0.16, p = 0.001$) relative to the reference variable (response: incorrect action). Delayed trigger failures are defined as failures in which the voice assistant experiences latency when activating, to the point of potentially, but not necessarily, providing a correct response too late to be useful. They had no measurable effect on benevolence ($m = 55.6, \beta = 0.043, p = 0.306$). None of the participants reported a failure due to a delayed trigger in interviews.

6.2 Perception Failures

Users reported failures across all four failure types reported in Table 1, including truncation, overcapture, noisy channel, and transcription. Perception failures indicate that the voice assistant did not accurately capture the users’ input. Transcription was by far

	Ability				Benevolence				Integrity			
	betas	se	Z	p	betas	se	Z	p	betas	se	Z	p
(Intercept)	3.56	0.046	77.413	<0.001	3.796	0.048	79.083	<0.001	3.754	0.045	83.422	<0.001
General Trust	0.23	0.096	2.396	0.017	0.111	0.111	1	0.317	0.339	0.106	3.198	0.001
Missed Trigger	-0.10	0.048	-1.958	0.05	-0.228	0.043	-5.302	<0.001	0.194	0.037	5.243	<0.001
Spurious Trigger	0.36	0.046	7.761	<0.001	0.236	0.041	5.756	<0.001	0.212	0.037	5.73	<0.001
Delayed Trigger	0.16	0.046	3.391	0.001	0.043	0.042	1.024	0.306	0.249	0.036	6.917	<0.001
Truncation	0.16	0.046	3.522	<0.001	0.126	0.042	3	0.003	0.263	0.037	7.108	<0.001
Overcapture	-0.13	0.047	-2.809	0.005	-0.193	0.042	-4.595	<0.001	0.133	0.037	3.595	<0.001
Noisy Channel	0.14	0.046	2.978	0.003	0.076	0.042	1.81	0.07	0.326	0.036	9.056	<0.001
Transcription	-0.06	0.047	-1.213	0.225	-0.101	0.042	-2.405	0.016	0.093	0.037	2.514	0.012
No Understanding	0.06	0.047	1.17	0.242	-0.043	0.042	-1.024	0.306	0.28	0.037	7.568	<0.001
Misunderstanding	-0.017	0.046	-0.37	0.711	-0.023	0.042	-0.548	0.584	0.202	0.037	5.459	<0.001
Ambiguity	0.46	0.046	9.913	<0.001	0.302	0.042	7.19	<0.001	0.361	0.036	10.028	<0.001
No Action	-0.003	0.047	-0.064	0.949	-0.09	0.042	-2.143	0.032	0.186	0.037	5.027	<0.001

Table 4: The results of three mixed-linear regression models, demonstrating how voice assistant failures impact users’ trust in voice assistants’ across ability, benevolence, and integrity. Reference failure source: Incorrect Action.

the most common failure source, contrasted with only one failure recorded per truncation, overcapture, and noisy channel.

Truncation failures indicate that the voice assistant stopped listening to input too early, and only acted on some of the user’s intended input. P12 reported that “*I use [a voice assistant] to send messages and stuff, and sometimes it would write the text for some of the words, but not all of the words. So it takes me longer than expected to send a message, because it will take a little bit of the words and not fully listen.*” They said, “*it’s aggravating, very annoying.*” Truncation failures had a favorable relative impact on perceptions of ability ($m = 48.7, \beta = 0.16, p < 0.001$) and benevolence ($m = 58.1, \beta = 0.126, p = 0.003$). As shown in Fig. 3, these maintained higher relative trust compared to other failures in perception.

Overcapture failures indicate that the voice assistant has listened beyond the point that a user has given their input. As P8 said, sometimes, “*it doesn’t know when to search for what I said and just keeps listening without taking action, even though it shows it is listening.*” They tried to make sense of this failure, saying “*I find that on different devices, the reaction time for it [is different].*” They said that, “*This is wasting my time. Which is only logically two to three minutes,*” but they said, “*if you keep messing with it, it makes it worse.*” Failures due to overcapture were particularly harmful to users’ confidence in voice assistants’ ability ($m = 38.5, \beta = -0.10, p = .05$) and benevolence ($m = 47.0, \beta = -0.228, p < 0.001$), with the overall lowest means compared to all other failure types.

There was one instance in which a user thought that the failure they experienced was because of noise in the background, indicative of noisy channel failures. P9 said, “*Sometimes...I’ll try to use a feature where it tries to identify like a song...and it just won’t be able to pick it up, and it’ll just give me a message, like ‘Sorry, I could not understand that.’*” They said, “*I get that it was loud...I would think that it would, it should be able to understand. So I feel like that is a little annoying.*” However, they said the failure did not impact how they thought about the voice assistant’s accuracy or ability, saying that “*it’s pretty accurate for the most part, for other things.*” Noisy channel failures were considered to more favorably impact user perceptions of ability ($m = 49.2, \beta = 0.14, p = 0.003$), with no measurable impact on benevolence ($m = 57.9, \beta = 0.076, p = 0.07$).

As shown in Fig. 3, they achieved similar levels of trust as failures due to truncation.

Nine of our participants mentioned failures relating to transcription of their input, indicating that they did not believe the voice assistant accurately captured what they had said. These failures varied from not understanding the name of a musical group (P7), incorrectly transcribing a text message (P2), incorrectly transcribing a sequence of numbers (P4), not understanding angry, slurred, or mumbled speech (P3, P5, P9), and not understanding accents (P8) or other languages (P6, P9). P7 said it caused a “*tiny little bit of frustration*” when it did not understand the musician they were requesting. However, they “*don’t really demerit [the voice assistant] for that in particular because it’s so good at everything else that it does.*” However, when it came to using the voice assistants in other languages such as Spanish or French, “*there has not been a successful time where it’s been able to play that different song in a different language*” (P9). This led the participant to think, “*that it – it just has no ability to understand me in a different language*” (P9). Failures due to transcription did not have a measurable impact on perceptions of ability in the survey ($m = 42.3, \beta = -0.06, p = 0.225$) relative to the reference variable, however they impacted trust more so than other failure sources within perception as shown in Fig. 3. Transcription failures did negatively impact perceptions of benevolence ($m = 51.0, \beta = -0.101, p = 0.016$).

6.3 Understanding Failures

We found that participants submitted failures across all categories of understanding failures, as described below.

Failures due to no understanding resulted in a complete inability to map the input to an action or response. P6 said, “*I was trying to plan a vacation...It was my friend’s bachelorette party...And I was like, [Voice Assistant], where’s Lake Havasu? How far is it?...And she’s like, ‘Sorry, I didn’t understand what you’re saying.’*” This led P6 to question, “*Why do I even use you?*” However, they said that, “*for timers, it works really well.*” No understanding failures did not significantly impact trust relative to the reference variable, in terms of ability ($m = 45.3, \beta = 0.06, p = 0.24$) or benevolence ($m = 52.5, \beta = -0.043, p = 0.306$).

Misunderstanding failures occurred when the voice assistant mapped the user's input to an action that was partially, but not fully, accurate to their intent. For example, P4 explained that when they ask their voice assistant "to 'Take me home.' It usually directs me to my home, but on occasion, it shows me search results for the phrase 'Take me home.'" Similarly, P1 explained how when using a voice assistant for online shopping, sometimes it would "pull up the wrong item or, like, the wrong location." They said they felt "disappointed and frustrated." Misunderstanding failures did not measurably impact perceptions of ability ($m = 42.4, \beta = -0.017, p = 0.711$) or benevolence ($m = 52.9, \beta = -0.023, p = 0.584$) relative to the reference variable.

Failures due to ambiguity were situations in which one could see several reasonable interpretations of one's intent from the captured input, but the system failed to navigate the ambiguity. For example, P10 said, "I was trying to get to Pizza Hut and... it kept on telling me one in the nearby city instead of the one that's I believe like 10 minutes away from me. So I asked a couple of times, and then it didn't work, and that's when I just pulled out my phone and then just looked it up myself and left." They said that they were "a bit baffled, since normally, like when I ask [a voice assistant] for something, I get the response I would expect." As demonstrated in Fig. 3 and Table 4, failures due to ambiguity were more favorable to users' impressions of the voice assistant's ability ($m = 62.6, \beta = 0.46, p < 0.001$) and benevolence ($m = 67.9, \beta = 0.302, p < 0.001$). Overall, these failures maintained the highest level of user trust.

6.4 Response Failures

There were two possible types of response failures. These included incorrect action, in which the system gives information that is incorrect, or no action, in which a voice assistant fails to respond at all.

Incorrect action failures were times when the command seemed to be accurately understood, but the information provided in response was incorrect. For example, P1 said that sometimes they would use "the voice assistant to give me the best route to get to the location." While it would usually accurately respond to this command, sometimes, "it will give me a really like roundabout way, like really time-consuming way." As shown in Fig. 3, failures due to incorrect action resulted in a relatively average perception of ability ($m = 44.0$) and benevolence ($m = 54.3$), and the lowest perception of integrity ($m = 53.6$).

Multiple users experienced failures due to no action, in which the voice assistant completely fails to respond to the input. P2 said, "I did have a couple times that was also frustrating... I would say 'Reply' [to a text message]. And I would talk and nothing would get sent. And like, my hands are literally covered in stuff because I'm rolling these cookies out, and I had to stop what I'm doing, go back to my phone, and actually like manually text." Another participant experienced failures due to no action, saying that "This morning where I woke up. I said, [Voice Assistant], what's the weather outside? And it loaded for the first few seconds... and then after a couple of seconds, it said, 'There was an error. Please try again in a few minutes.' I wait one or two seconds, then I'll ask it again, and it gives me the information" (P10). This participant said that because the information has been "accurate," they "would still trust it to a very high degree." Failures

due to no action had no measurable relative impact on ability ($m = 43.2, \beta = -0.003, p = 0.949$) and had a slight but significant negative impact on benevolence ($m = 50.7, \beta = -0.09, p = 0.032$) compared to incorrect action.

7 RESPONSES TO FAILURES AND FUTURE USE OF VOICE ASSISTANTS

Users described a variety of strategies for mitigating failures, given that they did occur. In some cases, users described completely stopping their use of a voice assistant for a particular task. For example, after encountering a truncation failure while using the voice assistant to send a text message, P12 said that they either "have to redo it, or I just, like, don't do it at all." Eventually, P12 said that they stopped encountering that failure because they "barely use it" for that same task anymore. So while some users felt like they "don't sweat it too much" (P5) when a voice assistant failed at a task, others felt like they would use it "not as much" (P2) for those same tasks.

We found that the pattern of continuing to use a voice assistant in general but excluding the tasks that resulted in a failure, at least for a short period of time, was consistent across many different types of failures, including transcription, misunderstanding, and ambiguity. For example, P2 said that they needed to be careful using a voice assistant, because sometimes they would say a name and "it would come up [with] a different name." They said that following an incident like that,

"I would still use [the voice assistant]. I think what would happen though is like you kind of build up that trust... So the next couple times I would go into my contacts and hit the button myself, you know, and then like if I was walking to my car and get my keys in one hand, and it's been a while. So, you know, let me try this again. Like I think that's something where you kind of have to like, build the trust back up and give it another try. At least that's what I do."

P12 echoed this, saying, "Let's say you're opening Spotify or something like that. I think it will probably go on command, rather than sending a message... different tasks, you know, it has a different trust level." P5 had a similar sentiment, saying "I think the problem with the most voice assistant is, if I tried to give it a complex search query, it doesn't really understand me, or it gets frustrating and I just I'm going to go ahead and type in whatever it is I'm looking for." Even when failures were mitigated in the moment, users remained wary of using their voice assistants for the same tasks.

Interestingly, sometimes users would continue to use their voice assistant for the same general task following a failure, but they would make slight changes to their use. For example, P1 encountered a misunderstanding failure while trying to shop for a sweater online, and they started to "rely on it a little less, and do more searching on my own." They said that "for future reference, I would just remember to not use it to do certain tasks and do certain tasks on my own, [especially] when I look for an item that's difficult to find." However, in the meantime, "I would just ask for other tasks." For P1, this included "looking for other items other than this sweater. I would tell her to search for like grocery items and do some comparison shopping online." Shopping for different items was distinct enough to maintain this user's trust. P7 experienced a similar situation, in

	Playing a Song				Texting a Coworker				Transferring Money			
	β	<i>se</i>	Z	<i>p</i>	β	<i>se</i>	Z	<i>p</i>	β	<i>se</i>	Z	<i>p</i>
Ability	0.048	0.003	16.000	<0.001	0.064	0.003	21.333	<0.001	0.076	0.005	15.2	<0.001
Benevolence	0.043	0.003	14.333	<0.001	0.028	0.003	9.333	<0.001	0.017	0.005	3.4	0.001
Integrity	0.019	0.002	9.500	<0.001	0.024	0.003	8	<0.001	0.03	0.004	7.5	<0.001
General Trust	-0.146	0.108	-1.352	0.176	-0.094	0.134	-0.701	0.483	-0.192	0.218	-0.881	0.378

Table 5: The results of three mixed-ordinal regressions modeling user trust in the voice assistant to execute the task based on their perceptions of the voice assistant’s ability, benevolence, and integrity. We did not include cut point calculations and state 1 calculations in the table for ease of interpretability.

which they encountered a transcription error, which they mitigated by spelling the name of “*hyerpop duo 100 gees*” as “*G-E-C-S*.” They said this correct helped so that “[*the assistant did*] understand what I was saying.” Even though they had experienced a failure for that particular artist, they “*continue to do that [use it to play songs] to this day. It’s a very good music player,*” but they are “*a little weary when it comes to certain musicians that I feel that... [the voice assistant] would have trouble understanding.*”

Users often made sense of the failures based on the perceived task complexity. P12 thought that the task that they had the highest trust in was “*to open like apps,*” followed by “*calling somewhere.*” They explained that, “*I want to put that as number one, but sometimes, like the way the contact name is, is not registered. Like, you know the way for you to say it, it’s not how like the voice [assistant] says it.*” P2 similarly evaluated the voice assistant, saying, “*the best thing is picking up website information.*” However, they similarly said “*to get more personalized messages, contacts, and that sort of thing, you have to be really careful what you say and how you say it.*”

Because of these findings, we hypothesized that users’ trust in voice assistants after failures would affect their willingness to use it for different tasks to differing degrees. As shown in Table 5, we used three mixed-ordinal regressions to model trust in these three tasks, with scores for confidence in the voice assistant’s ability, benevolence, and integrity as the independent variables. Trust in the voice assistant to play a song, text a coworker, and transfer money was encoded as an ordinal value. Confidence in the voice assistant’s ability, benevolence, and integrity were encoded as numerical values. General trust tendency was encoded as a numerical value and PID was encoded as a random categorical value. We found that user perceptions of voice assistant ability, benevolence, and integrity positively correlated with their willingness to use the voice assistant for future tasks. Overall, people were moderately trusting of their voice assistant to play a song ($m = 3.29, sd = 1.30$), less trusting of their voice assistant to text a coworker ($m = 2.34, sd = 1.18$), and least trusting of their voice assistant to transfer money ($m = 1.56, sd = 0.91$).

In particular, perceptions of ability had a stronger effect on people’s willingness to use the voice assistant to play a song ($\beta = 0.048, p < .001$) compared with benevolence, which also significantly impacted willingness to use the voice assistant to play songs, but to a slightly lesser degree ($\beta = 0.043, p < .001$). Integrity was even less influential, though still significantly positively correlated without how much people trusted their voice assistant to play a song ($\beta = 0.019, p < .001$). This pattern was repeated for texting a coworker and transferring money as well, with ability being most strongly positively correlated with people’s willingness to trust the

voice assistant to execute these tasks, followed by benevolence, and then integrity.

8 DISCUSSION

With interviews, a survey, and a crowdsourced voice assistant failures dataset, we conducted a mixed-method study of voice assistant failures and how they impact user trust in and future intended use of voice assistants. As the underlying technology for voice assistants continues to improve in accuracy and ability, and its applications become increasingly high stakes to human health and well-being [17, 41, 55, 66], we discuss our findings with the goal of improving user trust and long-term engagement in voice assistants.

Our users consistently relied on their voice assistants to find information and execute tasks across varying levels of complexity. Similar to prior work [35], those who wanted to use a voice assistant consistently for tasks which might result in failures have developed complex mental models of which tasks they can trust their voice assistants with. Unlike prior work [35], people often did not necessarily entirely abandon the use of their voice assistant after it failed at complex tasks, even after repeated failures. Many users considered the accuracy of their voice assistants so consistently high that they could forgive failures and continue engaging with those tasks after a short period of time. While trust in the complex tasks was being repaired, many participants continued using their voice assistants for tasks they considered more simple, such as information retrieval and playing music.

We find that failures that lead users to feel like they have wasted time, such as those due to missed triggers and overcapture, tend to lead to more deteriorated perceptions of ability and benevolence. This is contrasted with scenarios in which users have more understanding of *why* the voice assistant failed, such as those due to ambiguity and transcription, which users generally felt like they could work around or anticipate. However, if the failure due to transcription was believed to be due to using the device in another language, this caused abandonment of the voice assistant in that language. Similarly, users did not feel like they lost out on the advantages of using voice assistants when spurious trigger failures occurred, so they were less damaging to perceptions of ability. The single most damaging failure source to voice assistant integrity was action execution: incorrect action, as participants were more skeptical of the claim that the voice assistant would not cause harm following these failures.

Prior work has pointed to ways that trust can be repaired when failures do occur. Cuadra et al. [16] showed that when a voice assistant proactively attempts to acknowledge a failure and repair trust,

this increased people’s perception of its intelligence. Additionally, Mahmood et al. [39] has found that failure mitigation strategies such as apologies were effective in restoring perceptions of likability and intelligence of a voice assistant after a failure. Xiao et al. [65] demonstrated that situating the voice assistant as a learner, and helping users understand when to give feedback to the voice assistant, improved users’ perceptions of the voice assistant. Fischer et al. [23] encourages voice assistant responses to support progressivity of the conversation, especially when the response does not help the user. Our work shows that users naturally repair trust with their voice assistants by relying on it for different tasks following a failure, or the same task but on a different topic, such as online shopping for different items or playing music by other artists than those that caused a failure.

Quantitatively, we established that certain types of failures are more critical than others. This insight can be used to help prioritize the failure recovery strategies across HCI and NLP that are most effective for regaining trust. For example, self-repair for voice assistants such as Cuadra et al. [16] employed may be most useful in situations where the voice assistant has failed because of a missed trigger or overcapturing users’ input. In addition, we can also try to identify the specific components in the voice assistant technology stack that cause critical failures, and leverage techniques in NLP robustness to improve how these models perform during user interactions. For example, noisy channel and transcription failures can be modeled as small perturbations to the input, which is well researched [4, 21]. Reliable transcription may also be important to address by speech recognition modules, especially for low resource languages [38].

Our open-sourced dataset has also provided concrete and comprehensive example failures (199 real-world sourced examples with context, query, and response) for future researcher to reuse to develop failure mitigation strategies, along with a refined taxonomy for classifying voice assistant failures, as supported by prior work. While prior work [28] was useful in helping NLP practitioners anticipate and plan for failures across many types of NLP technologies, our dataset specifically addresses failures that occur with voice assistants. We anticipate this will allow future researchers to use human-centered example failures when conducting research related to voice assistant failures, trust, and mitigation strategies.

LIMITATIONS

There are a few methodological limitations of our study, which we detail here. First, our dataset collection and interviews relied on retrospectives and recalling failures, rather than observing them *in situ*. This subjects our data to recall bias, and our results should be interpreted in this light. For instance, none of the interview participants recalled failures due to missed triggers or delayed triggers in interviews, although missed triggers were considered relatively damaging to perceptions of ability and benevolence in the survey. Additionally, our survey relied on collecting users’ feedback regarding hypothetical scenarios. Future work may build on our findings by using our dataset to systematically introduce failures and capture the resulting impact on user trust via ESM or diary study. Our sample of participants were also frequent voice assistant users, which indicates that they likely forgave errors more easily

than other populations [35]. Additionally, we did not address the use of conversational agents through interfaces other than voice, such as embodied conversational agents or text-based conversational agents. As embodied and text interfaces have more potential affordances with which users can judge and interact with the system [5, 6], the impact of failures may not perfectly generalize to these use cases.

9 CONCLUSION

In conclusion, through a mixed-method study, we found that voice assistant users experience a multitude of failures, ranging from a voice assistant incorrectly triggering to responding in a way that does not address users’ needs. These different types of failures do differentially impact users’ trust, which in turn affects intention to use their voice assistants for tasks in the future. In particular, we find that failures due to spurious triggers and ambiguity are less detrimental to user trust than failures due to incorrect action execution, missed triggers, or overcapture. We additionally find that people rebuild their trust in voice assistants through simple tasks, such as playing a song, before resuming using their full voice assistant functionality after a failure has occurred. We also contribute a dataset of 199 failures, to help future researchers and practitioners build on our work. By further working to understand, prevent, and repair voice assistant failures, we hope to build voice assistant users’ trust in these devices and allow them to benefit from the increasing and varied functionality they provide.

REFERENCES

- [1] Deepali Aneja, Daniel McDuff, and Mary Czerwinski. 2020. Conversational Error Analysis in Human-Agent Interaction. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*. Association for Computing Machinery, New York, NY, USA, Article 3, 8 pages. <https://doi.org/10.1145/3383652.3423901>
- [2] Ron Artstein, Jill Boberg, Alesia Gainer, Jonathan Gratch, Emmanuel Johnson, Anton Leuski, Gale Lucas, and David Traum. 2018. The Niki and Julie Corpus: Collaborative Multimodal Dialogues between Humans, Robots, and Virtual Agents. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA), Miyazaki, Japan. <https://aclanthology.org/L18-1463>
- [3] Tammy Bahmanziari, J Michael Pearson, and Leon Crosby. 2003. Is trust important in technology adoption? A policy capturing approach. *Journal of Computer Information Systems* 43, 4 (2003), 46–54.
- [4] Yonatan Belinkov and Yonatan Bisk. 2017. Synthetic and natural noise both break neural machine translation. *arXiv preprint arXiv:1711.02173* (2017).
- [5] Timothy Bickmore and Justine Cassell. 2001. Relational Agents: A Model and Implementation of Building User Trust. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Seattle, Washington, USA) (CHI '01). Association for Computing Machinery, New York, NY, USA, 396–403. <https://doi.org/10.1145/365024.365304>
- [6] Timothy Bickmore and Justine Cassell. 2005. *Social Dialogues with Embodied Conversational Agents*. Springer Netherlands, Dordrecht, 23–54. https://doi.org/10.1007/1-4020-3933-6_2
- [7] Michael Braun, Anja Mainz, Ronnee Chadowitz, Bastian Pflöging, and Florian Alt. 2019. At Your Service: Designing Voice Assistant Personalities to Improve Automotive User Interfaces. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300270>
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [9] Heloisa Candello, Claudio Pinhanes, Mauro Pichiliani, Paulo Cavalin, Flavio Figueiredo, Marisa Vasconcelos, and Haylla Do Carmo. 2019. The Effect of Audiences on the User Experience with Conversational Interfaces in Physical Spaces. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3290605.3300320>
- [10] Herbert H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/CBO9780511620539>
- [11] Leigh Clark, Nadia Pantidi, Orla Cooney, Philip Doyle, Diego Garaialde, Justin Edwards, Brendan Spillane, Emer Gilmartin, Christine Murad, Cosmin Munteanu, Vincent Wade, and Benjamin R. Cowan. 2019. What Makes a Good Conversation? Challenges in Designing Truly Conversational Agents. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300705>
- [12] Jamie Condliffe. 2017. AI Voice Assistant Apps are Proliferating, but People Don't Use Them. MIT Technology Review. <https://www.technologyreview.com/2017/01/23/154449/ai-voice-assistant-apps-are-proliferating-but-people-dont-use-them/>
- [13] Toby A. Cox. 2020. Siri and Alexa Fails: Frustrations With Voice Search. <https://themanifest.com/digital-marketing/resources/siri-alex-fails-frustrations-with-voice-search>
- [14] John W Creswell and Cheryl N Poth. 2016. *Qualitative inquiry and research design: Choosing among five approaches*. Sage publications, Thousand Oaks.
- [15] Andrea Cuadra, Hansol Lee, Jason Cho, and Wendy Ju. 2021. Look at Me When I Talk to You: A Video Dataset to Enable Voice Assistants to Recognize Errors. *arXiv preprint arXiv:2104.07153* (2021).
- [16] Andrea Cuadra, Shuran Li, Hansol Lee, Jason Cho, and Wendy Ju. 2021. My Bad! Repairing Intelligent Voice Assistant Errors Improves Interaction. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–24.
- [17] Celso M De Melo, Kangsoo Kim, Nahal Norouzi, Gerd Bruder, and Gregory Welch. 2020. Reducing Cognitive Load and Improving Warfighter Problem Solving With Intelligent Virtual Assistants. *Frontiers in psychology* 11 (2020), 554706.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [19] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144, 1 (2015), 114.
- [20] Mary T Dzindolet, Linda G Pierce, Hall P Beck, and Lloyd A Dawe. 2002. The perceived utility of human and automated aids in a visual detection task. *Human factors* 44, 1 (2002), 79–94.
- [21] Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. HotFlip: White-Box Adversarial Examples for Text Classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Melbourne, Australia, 31–36. <https://doi.org/10.18653/v1/P18-2006>
- [22] Yuksel Ekinci and Philip L Dawes. 2009. Consumer perceptions of frontline service employee personality traits, interaction quality, and consumer satisfaction. *The service industries journal* 29, 4 (2009), 503–521.
- [23] Joel E. Fischer, Stuart Reeves, Martin Porcheron, and Rein Ove Sikveland. 2019. Progressivity for Voice Interface Design. In *Proceedings of the 1st International Conference on Conversational User Interfaces* (Dublin, Ireland) (CUI '19). Association for Computing Machinery, New York, NY, USA, Article 26, 8 pages. <https://doi.org/10.1145/3342775.3342788>
- [24] Jack FitzGerald, Shankar Ananthakrishnan, Konstantine Arkoudas, Davide Bernardi, Abhishek Bhagia, Claudio Delli Bovi, Jin Cao, Rakesh Chada, Amit Chauhan, Luoxin Chen, Anurag Dwarakanath, Satyam Dwivedi, Turan Gojayev, Karthik Gopalakrishnan, Thomas Gueudr, Dilek Hakkani-Tur, Wael Hamza, Jonathan J. Hüser, Kevin Martin Jose, Haidar Khan, Beiye Liu, Jianhua Lu, Alessandro Manzotti, Pradeep Natarajan, Karolina Owczarzak, Gokmen Oz, Enrico Palumbo, Charith Peris, Chandana Satya Prakash, Stephen Rawls, Andy Rosenbaum, Anjali Shenoy, Saleh Soltan, Mukund Harakere Sridhar, Lizhen Tan, Fabian Triefenbach, Pan Wei, Haiyang Yu, Shuai Zheng, Gokhan Tur, and Prem Natarajan. 2022. Alexa Teacher Model: Pretraining and Distilling Multi-Billion-Parameter Encoders for Natural Language Understanding Systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (Washington DC, USA) (KDD '22). Association for Computing Machinery, New York, NY, USA, 2893–2902. <https://doi.org/10.1145/3534678.3539173>
- [25] Karthik Gopalakrishnan, Behnam Hedayatnia, Qinlang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. Topical-Chat: Towards Knowledge-Grounded Open-Domain Conversations. In *Proc. Interspeech 2019*. 1891–1895. <https://doi.org/10.21437/Interspeech.2019-3079>
- [26] Aditya Gupta, Jiacheng Xu, Shyam Upadhyay, Diyi Yang, and Manaal Faruqui. 2021. Disfl-QA: A Benchmark Dataset for Understanding Disfluencies in Question Answering. In *Findings of ACL*.
- [27] Paige Hayes, Jason Wagner, Mark McCaffrey, and Matt Hobbs. 2018. Consumer intelligence series: Prepare for the Voice Revolution. <https://www.pwc.com/us/en/services/consulting/library/consumer-intelligence-series/voice-assistants.html>
- [28] Matthew K. Hong, Adam Fournier, Derek DeBellis, and Saleema Amershi. 2021. Planning for Natural Language Failures with the AI Playbook. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 386, 11 pages. <https://doi.org/10.1145/3411764.3445735>
- [29] Tatum Hunter. 2022. Siri and Alexa are getting on their owners' last nerves. The Washington Post. <https://www.washingtonpost.com/technology/2022/03/07/voice-assistants-wrong-answers/>. (Accessed on 06/08/2022).
- [30] Rinat Khaziev, Usman Shahid, Tobias Röding, Rakesh Chada, Emir Kapanci, and Pradeep Natarajan. 2022. FPI: Failure Point Isolation in Large-scale Conversational Assistants. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*. Association for Computational Linguistics, Hybrid: Seattle, Washington + Online, 141–148. <https://doi.org/10.18653/v1/2022.naacl-industry.17>
- [31] Dounia Lahoual and Myriam Frejus. 2019. When Users Assist the Voice Assistants: From Supervision to Failure Resolution. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290607.3299053>
- [32] Chia-Hsuan Lee, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee. 2018. Spoken SQuAD: A Study of Mitigating the Impact of Speech Recognition Errors on Listening Comprehension. *Proc. Interspeech 2018* (2018), 3459–3463.
- [33] One-Ki Daniel Lee, Ramakrishna Ayyagari, Farzaneh Nasirian, and Mohsen Ahmadian. 2021. Role of interaction quality and trust in use of AI-based voice-assistant systems. *Journal of Systems and Information Technology* 23, 2 (2021), 154–170.
- [34] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [35] Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA": The Gulf between User Expectation and Experience of Conversational Agents. 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, 5286–5297. <https://doi.org/10.1145/2858036.2858288>
- [36] Xiao Ma, Jeffrey T. Hancock, Kenneth Lim Mingjie, and Mor Naaman. 2017. Self-Disclosure and Perceived Trustworthiness of Airbnb Host Profiles. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (Portland, Oregon, USA) (CSCW '17). Association for Computing Machinery, New York, NY, USA, 2397–2409. <https://doi.org/10.1145/2998181.2998269>

- [37] Maria Madsen and Shirley Gregor. 2000. Measuring human-computer trust. In *11th Australasian conference on information systems*, Vol. 53, 6–8.
- [38] Alexandre Magueresse, Vincent Carles, and Evan Heetderks. 2020. Low-resource languages: A review of past work and future challenges. *arXiv preprint arXiv:2006.07264* (2020).
- [39] Amama Mahmood, Jeanie W Fung, Isabel Won, and Chien-Ming Huang. 2022. Owing Mistakes Sincerely: Strategies for Mitigating AI Errors. 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 578, 11 pages. <https://doi.org/10.1145/3491102.3517565>
- [40] Roger C Mayer, James H Davis, and F David Schoorman. 1995. An integrative model of organizational trust. *Academy of management review* 20, 3 (1995), 709–734.
- [41] Nikita Mehandru, Samantha Robertson, and Niloufar Salehi. 2022. Reliable and Safe Use of Machine Translation in Medical Settings. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 2016–2025. <https://doi.org/10.1145/3531146.3533244>
- [42] John Miller, Karl Krauth, Benjamin Recht, and Ludwig Schmidt. 2020. The Effect of Natural Distribution Shift on Question Answering Models. In *Proceedings of the 37th International Conference on Machine Learning (ICML '20)*. JMLR.org, Article 641, 12 pages.
- [43] Chelsea Myers, Anushay Furqan, Jessica Nebolsky, Karina Caro, and Jichen Zhu. 2018. Patterns for How Users Overcome Obstacles in Voice User Interfaces. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (CHI '18). Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3173574.3173580>
- [44] Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. *arXiv preprint arXiv:1806.00692* (2018).
- [45] Farzaneh Nasirian, Mohsen Ahmadian, and One-Ki Daniel Lee. 2017. AI-Based Voice Assistant Systems: Evaluating from the Interaction and Trust Perspectives. In *Americas Conference on Information Systems*.
- [46] Tim Paek and Eric J Horvitz. 2013. Conversation as action under uncertainty. *arXiv preprint arXiv:1301.3883* (2013).
- [47] Yury Pinsky. 2021. Loud and clear: AI is improving Assistant conversations. <https://blog.google/products/assistant/loud-and-clear-ai-improving-assistant-conversations/>. (Accessed on 12/12/2022).
- [48] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for SQuAD. *arXiv preprint arXiv:1806.03822* (2018).
- [49] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250* (2016).
- [50] Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics* 7 (2019), 249–266.
- [51] Samantha Robertson and Mark Diaz. 2022. Understanding and Being Understood: User Strategies for Identifying and Recovering From Mistranslations in Machine Translation-Mediated Chat. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 2223–2238. <https://doi.org/10.1145/3531146.3534638>
- [52] Tulika Saha, Saichethan Reddy, Anindya Das, Sriparna Saha, and Pushpak Bhat-acharyya. 2022. A Shoulder to Cry on: Towards A Motivational Virtual Assistant for Assuaging Mental Agony. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 2436–2449. <https://doi.org/10.18653/v1/2022.naacl-main.174>
- [53] Maha Salem, Gabriella Lakatos, Farshid Amirabdollahian, and Kerstin Dautenhahn. 2015. Would You Trust a (Faulty) Robot? Effects of Error, Task Type and Personality on Human-Robot Cooperation and Trust. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction* (Portland, Oregon, USA) (HRI '15). Association for Computing Machinery, New York, NY, USA, 141–148. <https://doi.org/10.1145/2696454.2696497>
- [54] William Seymour and Max Van Kleek. 2021. Exploring Interactions Between Trust, Anthropomorphism, and Relationship Development in Voice Assistants. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–16.
- [55] Emre Sezgin, Yungui Huang, Ujjwal Ramtekkar, and Simon Lin. 2020. Readiness for voice assistants to support healthcare delivery during a health crisis and pandemic. *NPJ Digital Medicine* 3, 1 (2020), 1–4.
- [56] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990* (2022).
- [57] Yueming Sun and Yi Zhang. 2018. Conversational Recommender System. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). Association for Computing Machinery, New York, NY, USA, 235–244. <https://doi.org/10.1145/3209978.3210002>
- [58] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. LaMDA: Language Models for Dialog Applications. *arXiv preprint arXiv:2201.08239* (2022). <https://arxiv.org/abs/2201.08239>
- [59] Suzanne Tolmeijer, Naim Zierau, Andreas Janson, Jalil Sebastian Wahdatehagh, Jan Marco Marco Leimeister, and Abraham Bernstein. 2021. Female by Default? – Exploring the Effect of Voice Assistant Gender and Pitch on Trait and Trust Attribution. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (CHI EA '21). Association for Computing Machinery, New York, NY, USA, Article 455, 7 pages. <https://doi.org/10.1145/3411763.3451623>
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [61] Julia Velkovska and Moustafa Zouinar. 2018. The illusion of natural conversation: interacting with smart assistants in home settings. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*.
- [62] Xuezhi Wang, Haohan Wang, and Diyi Yang. 2021. Measure and Improve Robustness in NLP Models: A Survey. *arXiv preprint arXiv:2112.08313* (2021).
- [63] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652* (2021).
- [64] Peter Winter and Annamaria Carusi. 2022. 'If You're Going to Trust the Machine, Then That Trust Has Got to Be Based on Something': Validation and the Co-Constitution of Trust in Developing Artificial Intelligence (AI) for the Early Diagnosis of Pulmonary Hypertension (PH). *Science & Technology Studies* 35, 4 (2022), 58–77.
- [65] Ziang Xiao, Sarah Mennicken, Bernd Huber, Adam Shonkoff, and Jennifer Thom. 2021. Let Me Ask You This: How Can a Voice Assistant Elicit Explicit User Feedback? *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 388 (oct 2021), 24 pages. <https://doi.org/10.1145/3479532>
- [66] Samuel Yang, Jennifer Lee, Emre Sezgin, Jeffrey Bridge, Simon Lin, et al. 2021. Clinical advice by voice assistants on postpartum depression: cross-sectional investigation using Apple Siri, Amazon Alexa, Google Assistant, and Microsoft Cortana. *JMIR mHealth and uHealth* 9, 1 (2021), e24045.
- [67] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300509>
- [68] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? *arXiv preprint arXiv:1801.07243* (2018).