Artificial Interactions: The Ethics of Virtual Assistants

by

Julia Guy

A thesis submitted in partial fulfillment of the requirements for the degrees of

Master of Arts in Digital Humanities

and

Master of Library and Information Studies

Digital Humanities and School of Library and Information Studies

University of Alberta

Abstract

Using an online survey, this research explores public opinion related to ethical concerns and the portrayal of human-like characteristics by Virtual Assistants (VAs), such as Apple's Siri and Amazon's Alexa. This research is centered around two research questions: 1) What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?; and, 2) What can be determined by evaluating public opinion about if and how Virtual Assistants should perform human-like characteristics and behaviours? Results of the survey were summarized and analyzed using a descriptive statistics approach. Findings indicate that there are multiple issues with VA technologies that produce significant levels of concern for respondents. Concerns related to corporate manipulation, surveillance and lack of transparency, as well as concerns related to political manipulation or increased political polarization, were the issues where participants most frequently expressed high levels of concern. Responses related to the portrayal of human-like characteristics by VAs suggest that, in many ways, public opinion might align with the way these technologies are currently designed, particularly in terms of what personality traits VAs should portray. There were, however, several instances where respondent opinions sharply contrasted the ways that VAs are currently being designed, particularly when it comes to what gender VAs should portray, and how human-like these technologies should be.

Preface

This thesis is an original work by Julia Guy. No part of this thesis has been previously published. This research project recieved research ethics approval from the University of Alberta Research Ethics Board, Artificial Interactions: The Ethics of Virtual Assistants, No. Pro00110059, July 4th 2021.

Dedication

For my parents, Fred and Mary Guy, and my brother, Tom.

Acknowledgments

I'd like to thank my excellent supervisors, Dr. Tami Oliphant and Dr. Geoffrey Rockwell, for their encouragement and thought-provoking feedback throughout the writing of this thesis. They are both exceptional academics and wonderful people and I am very fortunate to have had the chance to learn from them.

Thanks to Dr. Danielle Allard, my SLIS academic advisor, who was a terrific sounding board as I waded through ideas for this thesis.

I would also like to acknowledge the ongoing support of Dr. Michael McNally, who has been my greatest champion throughout my graduate experience. I know more than I ever wanted to about copyright and I have him to thank. My advice to anyone entering a masters program will forever be, "Find yourself a Michael McNally." Although I think there's only one out there.

I would also like to thank my fellow Digital Humanities students, in particular Anna Solazzo and Schyler Palm, for their encouragement and helpful ideas at our weekly thesis-group meetings.

Thank you to Rylan Kafara for his excellent research advice, big heart, and constant support.

Shout out to desk #5-113 in Rutherford North. You are the best desk on campus.

This work has been generously supported by the Social Sciences and Humanities Research Council (Master's SSHRC).

Contents

1. Introduction ……..….. 1

2. Literature Review……..……. 3

3. Methods……. 17

4. Respondent Demographics and Virtual Assistant Use……..... 25

5. Ethical Concerns about Virtual Assistants……..….. 37

6. Opinions on the Performance of Humanlike Characteristics…… 49

7. Conclusion…… 68

References …… 73

Appendices…… 84

List of Figures

Figure 1: *Age distribution of participants compared to the 2016 population of Alberta (2016b).*…… 27

Figure 2: *Levels of Concern for "Manipulation by companies" and "Threats to autonomy (the ability to make your own decisions)".*.....…..39

Figure 3: *Respondent Trust that Companies Ensure Technologies Do Not Have Negative Impacts*…….. 45

Figure 4: *Opinions on Who Should be Responsible for Preventing Negative Impacts of VAs*…… 47

Figure 5: *Responses to the Question, "How human-like should VAs be?"*……... 50

Figure 6: *Response to the Question, "What gender should a VA portray?"*……... 54

Figure 7: *Gender of Participants*…….. 55

Figure 8: *Percentage of Respondents who Selected Each Task-related Trait*....……. 59

Figure 9: *Percentage of Respondents who Selected Each Personality-related Trait*…… 60

List of Tables

Table 1: *How Participants Describe Using VA Technology*…… 33

Table 2: *The Ten VA Concerns Respondents Were Most Frequently "Very" or "Extremely" Concerned About*…… 37

Table B1: *Percentage of the 2016 Alberta Population for Each Age Group Option Provided to Respondents* ……101

Table D1: *Full List of VA Concerns, Rated in Order of Percentage of Respondents Who Were "Very" or "Extremely" About Each Issue* …… 105

Chapter 1: Introduction

Virtual Assistants (VAs) such as Apple's Siri and Amazon's Alexa, are examples of relation-based, or Social Artificial Intelligences (AIs). These systems are built around conversational AI technology and use natural language to respond to verbal questions and commands. Due to the rapid advancement of language-based AI technology, research exploring ethical concerns related to VAs and other socal AI technologies is urgently needed to ensure that the design and regulation of these technologies aligns with the opinions and priorities of the public. This research aims to reveal potential trends in public opinion around VAs, identify opportunities for future research in this area, and consider the ethical implications of VAs.

VAs have become popular around the world. This means that even though this technology is relatively new, huge portions of the population are using or at least familiar with VAs, and can therefore provide opinions on issues related to these technologies. The combination of urgency for further research on a quickly advancing technology and widespread familiarity with VAs, makes this an opportune time to gather public opinion data on this topic.

This research addresses the following questions:

1. What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?
2. What can be determined by evaluating public opinion about if and how Virtual Assistants should perform humanlike characteristics and behaviour?

This exploratory research used an online survey to collect opinion data. This data was collected through a self-selection process and uses non-probability, volunteer sampling. Because the data was gathered using convenience sampling, the results of the survey should not be interpreted as generalizable findings but are rather meant to provide indications of potential trends in public opinion around VAs that can reveal opportunities for future research.

Responses to the survey were analyzed using a descriptive statistics approach that involved: describing, organizing, and summarizing responses; using cross-tabs for further analysis; interpreting the results by conducting additional research; considering possible explanations for responses; and, identifying areas that require further research.

This thesis has been organized into seven chapters. Following this introduction is a chapter devoted to the literature review. The literature review will outline previous research done in several relevant areas including: surveillance capitalism, the anthropomorphization of

technology, the potential for manipulation by VAs, regulation and recommendations, public opinion research, and privacy. Chapter 3 discusses the research methods, including the process taken to develop the survey, and the approaches used to analyze the responses. The findings of the survey are then discussed in Chapters 4-6. The analysis begins with a chapter focussed on describing the respondents and their use of VAs, which provides a foundation for applying the results of the survey to the research questions. Chapter 5 focuses on the first research question by discussing the survey responses related to ethical concerns, such as those around trust, responsibility, privacy, and transparency. Chapter 6 addresses the second research question by interpreting the respondent's opinions on the portrayal of characteristics, such as age, gender, and personality traits by VAs. This chapter also discusses issues of AI self-disclosure, and how human-like the participants think VAs should be. The final chapter of this thesis restates key findings and summarizes the recommendations for future research that emerged throughout the analysis.

Chapter 2: Literature Review

The emergence and integration of Social AI has led to an increase in scholarly research around these technologies. This literature review begins with an introduction to Virtual Assistants (VAs), the Social AI technology which this research focuses on. This will be followed by an overview of surveillance capitalism, the rise of which laid the groundwork for the development and integration of VAs into daily life. This will be followed by an exploration of previous research in several relevant areas including: the anthropomorphization of technology, the potential for manipulation by AIs, privacy concerns, public opinion research, and recommendations or guidelines for the ethical development of AI. Although this research will focus primarily on ethical dilemmas relating to VAs it should be noted that these technologies are only one of many forms of Social AI and a small portion of the surveillance capitalism system.

Virtual Assistants

VAs, such as Apple's Siri and Amazon's Alexa, are examples of relation-based, or Social Artificial Intelligences (AIs). VAs respond to natural language commands and queries either by completing a task, providing a natural language response, or both. Built around conversational AI technology, VAs, sometimes called other names such as Voiced-based Digital Assistants (VBDAs), process complex data to facilitate information retrieval or the performance of tasks in real time (Vimalkumar et al., 2021). Conversational AIs, like VAs, are being designed to generate and communicate responses that accurately mimic human conversation.

All VAs are designed to respond verbally to voice cues and commands. There are also different design features and capabilities that vary depending on the VA. Alexa, for example, can be customized by adding "skills" which are downloaded from an online store that allow users to access countless apps through Alexa, including those connected to financial institutions, appliances, games, media outlets, and almost anything you could think of (Vigliarolo, 2020). Although VA systems vary in terms of capabilities and popularity, the ability to collect data and communicate verbally is common to all VAs.

History of VAs

VAs are a relatively new technology; however, scientists and computer engineers have been working toward voice-recognition and machine speech for a long time. In the 1700s people were already attempting to create machines that could mimic human speech, but significant progress began in the first half of the 20th century, led by Bell Laboratory scientists Harvery Fletcher and

Homer Dudley (Juang & Rabiner, 2005). By the 1950s, Bell laboratories had developed the "Audrey" system which recognized spoken numbers (Pinola, 2011). Voice recognition technology continued to advance and in the 1960s, William C. Dersch with IBM, developed a technology known as Shoebox, which could respond to voice commands by solving basic arithmetic problems (IBM, n.d.). The 1980s also saw significant advancement in voice-recognition technology due to the advancement of the "hidden Markov model" which incorporated statistics and probability so that technologies could guess words using probability (Pinola, 2011). Progress stalled somewhat before 2001, until Google was able to harness huge amounts of user data from its search engine to train voice recognition technology, which contributed greatly to accuracy (Pinola, 2011). The collection of vast amounts of data was therefore an essential step in improving the accuracy of voice recognition technology.

Apple's Siri was the first popular VA, first integrated into the iPhone 4 in 2011, the technology had originally been developed through decades of AI research through the Stanford Research Institute (SRI International, n.d.). In 2014 Amazon launched its VA as part of the smart speaker Amazon Echo (Etherington, 2014). Less than ten years later, Amazon's Alexa is available on hundreds of millions of devices from Amazon and third-party device manufacturers, according to an Amazon website (Amazon, n.d.c). The use of VAs has become mainstream and is growing worldwide (E-marketer, 2019). VAs are accessible most commonly through phones, smart speakers, and in the car. Today, VAs are "ubiquitously occupying our living spaces" (Vimalkumar et al., 2021).

On May 18th 2021, Google announced a new AI technology called Language Model for Dialogue Applications (LaMDA) that can chat with users about any subject (Hoa, 2021). LaMDA will eventually be integrated into Google's voice assistant and search portal, Workplace, as well as its collection of cloud-based work software including Gmail, Docs, and Drive. According to Goggle, LaMDA "can engage in a free-flowing way about a seemingly endless number of topics, an ability we think could unlock more natural ways of interacting with technology and entirely new categories of helpful applications" (Collins & Ghahramani, 2021).

VAs and Data Collection

Because these technologies gather data about us in our most intimate environments, like our homes or our cars, VAs are ideally positioned for maximum data collection. It is, however, difficult to determine exactly what data VAs collect and how it is used, both by the companies

that develop VAs and by third parties who may in some cases have access to data. The challenge of understanding how data is truly collected and used by VA companies is a moving target, as technology companies are changing rapidly and beginning to address privacy concerns. It is also likely that different companies such as Apple, Google, and Microsoft, use and collect user data in different ways. Certainly, one major use for data is to improve the technologies themselves (Amazon, n.d.b). Access to vast amounts of data has been the key to making voice-recognition technology's rapid advancement and growing accuracy (Pinola, 2011). Data is also used for personalization (Amazon, n.d.b; Apple n.d.) and VA developers continually promote the idea that by learning a user's patterns and interests makes these technologies more accurate and useful.

The fact that VAs are built for voice recognition inherently means that these systems can always be listening. Although Amazon insists that the system records nothing until a wake word, like "Alexa," is spoken (Amazon, n.d.b), the system can still record a lot of data within the queries or commands spoken to the device. At minimum, it seems that VAs currently record, store, and analyze voice commands spoken after the technology is prompted by the user. There are sometimes options for these recordings to be deleted by the user. For example, you can say specifically to Amazon's VA, "Alexa, delete everything I've ever said," and it will delete all of your voice recordings (Amazon, n.d.b). Although there appears to be features such as these, that allow for users to manage how their data is collected, it is up to the user to take their privacy into their own hands and find out how to delete their recordings themselves. Additionally, most of the information available about data collection and use is published by the companies themselves and is impossible for the average person to verify.

Apple's October 27 2021 Privacy Policy, vaguely states that "Apple uses personal data to power our services…" (Apple, 2021b). Statements like these do little to clarify how data is actually used, what is considered a service, etc.. Despite all of the unknowns about the specifics regarding what data is collected and how it is used by VA companies, what we know for certain is that these companies are profit-driven, have access to mass amounts of data, and, as will be discussed in further sections, are subject to very little regulation. These three factors create an opportunity for the exploitation of users and as such, ongoing research desperately needs to be conducted into the ethics of VAs.

Research on VAs

There are many ethical issues and social concerns around the use and design of VAs that have been explored previously, including: issues related to privacy (Sharif & Tenbergen, 2020); the potential negative impacts of "outsourcing" tasks to AI, such as increased ignorance and reduced responsibility for users (Frischmann & Selinger, 2018, Part II); and, issues related to preventing VAs from producing offensive responses (McTear, 2020, pp. 179-180).

There are also those who focus not on ethical issues with VAs, but rather their potential to *support* ethical decision-making and autonomy. Francisco Lara (2021) for example, argues that there is great potential for specially designed VAs to "teach users to make better moral decisions on their own" through dialogue designed to build user capacity to evaluate values and moral beliefs (Lara, 2021). Heinrich (2021) relatedly suggests that if standards for transparency are followed, VAs have the potential to support self-determination and self-improvement of users.

Notably, another significant portion of the research that has been done on VAs focuses on supporting the design of VA technology from a marketing perspective. Researchers have explored factors that influence consumer adoption of VAs (Vimalkumar et al., 2021; Foehr, J., & Germelmann, 2020), consumer willingness to make purchases through VAs (Christensen & Eriksen, 2021, Tassiello et al., 2021), and users willingness to disclose personal information to VAs (Ki et al., 2020). There does not appear to be significant research on VA portrayal of human-like characteristics which incorporate public opinion data, that does not come from a marketing perspective. This suggests that a lot of the research involving public opinion and VAs has been conducted to make these technologies more appealing to consumers, rather than to critically evaluate how these products are developed from the perspective of societal impacts or values.

With language based AI technology such as LaMDA rapidly advancing, there is great urgency to research public attitudes around these technologies. Researchers have a short window of time to gather and evaluate public opinion data before more sophisticated versions of VAs, equipped with technology such as LaMDA, become commonplace and entrenched in our devices, homes, and workplaces.

Surveillance Capitalism

Before investigating public opinion around VAs, it is essential to situate VAs as part of the larger capitalist technology landscape. VAs are primarily developed by powerful, profit-driven, technology companies, such as Google and Apple. VAs are designed to enhance

the convenience and appeal of the companies' technologies, while facilitating access to huge amounts of behavioural data.

Shoshana Zuboff (2019) defines surveillance capitalism as, "a new economic order that claims human experience as free raw material for hidden commercial practices of extraction, prediction, and sales...in which the production of goods and services is subordinated to a new global architecture of behavioural modification" (p. v). According to Zuboff, the surveillance of our online behaviour, or 'behavioural surplus,' began through tracking our use of search engines like Google and became the zero-cost asset that established the foundational mechanisms of surveillance capitalism (p. 81).

This surveillance by tech companies was originally presented as a service for the user (Zuboff, 2019, p. 69; Becker, 2019, p. 310). Companies argue that the more they know about you, the more personalized your experience will be. The data that is collected is often framed as helpful to the users, because it means that the ads they will be presented with will be more relevant and useful.

Contrary to popular belief, major players in the data collection world like Google and Facebook do not sell the data they collect, rather, they charge other companies for the *use* of people's data, without transferring the ownership of it (Christl, 2017, p. 11). They also permit third-parties to use their infrastructure to collect even more data, to benefit both the third-party client companies and the platforms themselves (p. 11). This system allows companies to maintain control and power over the behavioural data they collect.

The advancement of information technology has often made this collection and use of personal data invisible (Acquisti et al., 2015, p. 509). This means that it is difficult for the average person to conceive of how they are being tracked. The fact that technology companies' Terms and Conditions documents often do not say anything about the algorithms they use (Becker, 2019, p. 310) adds to the lack of understanding. Unfortunately, tech giants have skillfully exploited this lag in social evolution by outrunning public understanding and the development of law and regulation that this understanding produces (Zuboff, 2015, p.83). The swift and relentless rise of surveillance capitalism led to people unknowingly providing access to their behavioural data before the average person understood what was going on.

Lack of Regulation

Relinquishing control over personal data and privacy has shifted the balance of power from the subjects of the data to those holding the data (Acquisti, 2015, pp. 513-514). This massive power imbalance poses a huge challenge because there is so little regulation to reign-in big tech companies. This is because policy-making is slow moving by nature and governments cannot keep up with the rate at which technology is developing (Deloitte, 2018; Fenwick et al., 2017). Despite regulation being slow it still seems clear to some that the social problem of AI ethics will not be solved "by allowing markets to decide freely" (Whitby, 2012, p. 234) and that ethical interventions should be made.

As things currently stand, however, the ethical decisions that will affect our futures, democracy, and potentially even autonomy, seem to fall to the companies themselves. Engineers must use their own personal judgment and are trusted to discern and to evaluate the ethical stakes of their own products (Moss & Melcalf, 2019). It is unclear whether the relatively few individuals making these decisions can recognize the biases, long-term impacts, negative consequences, and potential for abuses, given their proximity to the technology. This is especially concerning considering that "the designers and operators of these technologies are often incentivized to subordinate others' interests to theirs" (Susser et al., 2019, p. 35) in an industry that relies on the collection of behavioural information and predicition.

Behavioural Prediction

Behavioural prediction is at the core of the surveillance capitalism business model which inevitably results in significant ethical issues. Whether it is predicting what people will do based on past behaviour, or based on what similar people (as defined by the technology) have done, "such systems may induce and extend the sort of conformity that they already presuppose" (King, 2020, p. 33). Predicting and monitoring has the power to shape what information is presented to us and thus our world view. Zuboff (2015) argues that this new business frontier goes beyond merely predicting our behaviour, and can now create opportunities to modify our actions for profit (p. 84).

Critics of Zuboff have argued that surveillance capitalism is not in fact a new economic system but rather an extension of what capitalism has always been about (Christophers, 2020). Regardless of whether surveillance capitalism is an entirely new economic order, or an extension of traditional capitalism, the concentration of power and ability to monitor, predict and influence consumer behaviour on a large scale, is unprecedented.

From Prediction to Manipulation

Scholars have noted the exceptionalism of this time and the shift from advertising that is based in persuasion and happens in plain sight, to more subtle manipulation by advertisers using personal data (Büchi et al., 2020, p. 7). Many people consider intrusive data collection as being only a minor problem, perhaps not appreciating how this shift can potentially lead to data influencing behavior at scale (Christl, 2017, pp. 5-6). Arguably, a consumer's free choice has always been essential in advertising ethics but "in the online world, this ethical value is scarcely met" (Becker, 2019, p. 310).

The relentless collection of data means that companies can move away from predictions, and the likely outcomes they provide, and toward guaranteed outcomes through manipulation. Susser et al. (2019) define online manipulation as "the use of information technology to covertly influence another person's decision- making" (p. 29). Zuboff (2019) outlines three main kinds of online manipulation: 1) tuning, which involves subliminal cues; 2) herding, which is done by controlling the situation and the options for response; and, 3) conditioning, which involves reinforcing positive behaviours, as determined by those designing the technology (p. 293-296).

Online manipulation may start in the digital sphere, but it does not stay there. Research suggests that messages that people receive digitally influence a variety of offline behaviours (Bond et al., 2012, p. 298). What people share can be used to influence their "emotions, thoughts, and behaviors in many aspects of their lives as individuals, consumers, and citizens" (Acquisti, 2015, p. 513-514). That said, it is not just the data we *share,* because the word "share" implies a conscious act, it is also the data about us and our behaviour that is collected in ways we may not even be aware of.

Arguably, this kind of manipulation can be especially effective because information technology fades into the background and out of people's awareness. "A determined manipulator could not dream up a better infrastructure through which to carry out his plans" (Susser et al., 2019, p. 24). This vulnerability to manipulation threatens our autonomy. This influence over our ability to think for ourselves could compromise society down the road given that "democratic institutions are designed (ideally) to reflect autonomous decisions reached in the political sphere" (p. 37). Acknowledging this potential, the EU High Level Expert Group on AI (2019) declared that "systems should not unjustifiably subordinate, coerce, deceive, manipulate, condition or herd humans" (p. 12).

Contrary to the above perspectives about the threat of manipulation, at least one study has predicted that AI in the form of robots could help promote a sense of autonomy, specifically for seniors living in assisted living environments (Pirhonen et al., 2020). That said, in most situations the sheer volume of personal data gathered and increasing potential for that data to be used for manipulation, combined with a lack of transparency by big tech companies, suggest that the threats to autonomy might outweigh the positive potential these technologies have to enhance it.

Anthropomorphization and Virtual Assistants

The surveillance capitalist business model hinges on the ability to predict and influence our behaviour. Now the goal for these companies is to find improved ways to gather increasingly personal data. Enter VAs.

A key factor that sets VAs apart from other popular information retrieval technologies like search engines, is that they are designed to communicate in the way that is most natural to humans, which is through conversation. When humans interact with each other verbally, our personalities and intentions are conveyed through our tone, inflection and the words we choose when conversing. By communicating with people through spoken language, VAs inevitably portray human characteristics. VAs have been programmed to respond with appropriate (if not yet entirely believable) language, tone, pitch, and inflection.

Although this communication method is very convenient, the anthropomorphization of technology can be problematic. Anthropomorphism being the act of "attributing humanlike capacities or properties to nonhuman agents" (Waytz et al., 2013). Studies have found that people tend to anthropomorphize Social AI, projecting our own mentality onto these systems (Foehr & Germelmann, 2020). The impersonation of human communication can lead to users endowing VAs and other Social AI technologies with human qualities and abilities that these technologies do not possess. For example, people could consciously or unconsciously assume VAs are built with qualities like ethical decision making or basic understandings of privacy, which humans are generally socialized to possess.

The creation of technology that can execute tasks and build relationships at the same time is a strategic move for surveillance capitalists, "who may mobilize this reliance to gain access to increasingly intimate types of information about their users" (Woods, 2018, p. 344). One study that examined user interactions with Alexa reported that users focussed more on the quality of

the interaction with the VA rather than whether Alexa had produced the desired end result (Lopatovska et al., 2019, p. 995). Similarly, another recent survey of Amazon Alexa users underscored that it is not just the functional capabilities "but also the emotional, relational, and social experience users have during their interactions" with their VAs that makes them successful (Ki et al., 2020).

Scheutz (2012) argues that the fact that people can develop attachments and unidirectional bonds with AIs, creates the potential for abuse and large-scale exploitation (pp. 211-217). The tendency to anthropomorphize VAs, combined with VAs' access to large amounts of personal data, creates great potential for persuasion or even manipulation to be executed through VA technology. It has been predicted that as anthropomorphic AI continues to improve, these systems will "surpass the most emotional persuasive technology we have seen" (Darling, 2017, p.179).

Many scholars have acknowledged the dangers of humanizing AI systems and offered various solutions. The Montréal Declaration for a Responsible Development of Artificial Intelligence stressed under its Principle of Autonomy, that AI should not be developed to create dependencies through attention-capturing techniques or the imitation of human characteristics in ways that could cause confusion between AIs and humans (Université de Montréal, 2018, p. 9). Tigard et al. (2020) suggested that they should be designed with social responsiveness and the ability to recognize emotions, but that they should not themselves be programmed to exhibit human emotions.

Others have proposed that there are great social benefits that could come from integrating human-like AI into our lives. For example, Tom Gruber (2017), one of the creators of Siri, argues that AI can be used to enhance our social lives by recording details from our social interactions so that we could retrieve peoples names and interests in the future. It is significant that VA's ability to collect personal data will only improve as companies continue to find ways to make these technologies more appealing and seemingly trustworthy.

Privacy Concerns

The ethical issues that come with anthropomorphization and the threat of manipulation discussed above, are intrinsically connected to the concept of privacy. Since online surveillance by corporations began, there have been public concerns about privacy. People's privacy concerns include worries over data collection and unauthorized sharing and manipulation with their

personal information (Anic et al., 2019, p. 801). Other concerns include things the leaking confidential information and government surveillance.

It has been argued that most privacy concerns on the part of the public are misguided, because in order for data to be considered 'personal information' it must identify and distinguish the person from everyone else in the universe, and the information that is collected and stored by computers is usually anonymized (Price, 2020, p.172). Indeed, tech companies and data brokers often use the argument that data is anonymized as a way to quell privacy concerns. However, this process of anonymization can more accurately be understood as 'pseudonymization' (Christl, 2017, p. 69). Users are given unique codes that can be matched to profiles and mapped to email addresses and phone numbers, which undermines claims of true 'anonymization' (p. 69). Although these companies may not know you by name (although they could probably find that out), your data can still be linked and used to predict and modify your behaviour.

On Apple's website dedicated to Siri it states, "Siri is designed to do as much learning as possible offline, right on your device. Searches and requests are not associated with your identity — so your personal information isn't gathered to sell to advertisers or other organizations" (Apple, n.d.). Notably, Apple reassures customers that their "personal information" will not be sold to other parties, but this (likely very carefully worded) statement makes it unclear whether this means your data is anonymized (or pseudonymized) and then sold to third parties, licensed to third parties rather than sold outright, or is not provided to other organizations in any form. It is also possible that even if third parties are not able to access this data it could still be used by Apple or one of its affiliates to influence behaviour or for other unexpected purposes.

The Guardian reported in 2019 that Apple contractors regularly hear confidential and personal information recorded through Siri as part of their quality-checking process and that it would not be difficult to identify the person that they are listening to (Hern, 2019). This practice of humans listening to recordings from VAs without explicit consent is one of the more overt examples of a violation of privacy. Although this example is potentially what people are likely to think of when they think of a privacy breach, there may also be more covert uses that people may not have considered.

Privacy and Public Opinion Research

There has been a great deal of research done on public opinion around privacy concerns and online surveillance. Shoshana Zuboff (2019) reports that of 48 prominent surveys on the

topic conducted between 2008 and 2017, 46 of them determined that substantial majorities support measures for enhanced privacy and user control over personal data (p.339). Despite the majority of people having concerns about privacy, the results of a 2019 study determined that many respondents hold a "relatively high level of fatalistic belief in technology and business' ability to protect privacy, and most respondents have fatalism towards the law" (Xie et al., 2019, p. 742). Furthermore, people with higher levels of fatalistic belief about technologies and business are less likely to protect their privacy on the internet (p. 742). The feeling that users should have more privacy and control over their data combined with understandable doubt in governments or companies ability to regulate tech companies, leads to an overarching sense of powerlessness.

This powerlessness is concerning given that an estimated 65% of Americans feel that major technology companies "often fail to anticipate how their products and services will impact society" (Smith, 2018, p. 3). A 2019 report on American public opinion around AI governance also found that Americans consider AI governance challenges to be important and that digital manipulation is a key issue that is likely to impact millions of people in the world in the next 10 years (Zhang et al., 2019, p. 17). The report also found that Americans with computer science engineering degrees, rate all AI governance challenges as less important than other members of the public (p. 20). This report did not discuss potential reasons why computer scientists might consider AI governance less concerning, it is possible this could be because they feel there is only minimal risk to the public based on their knowledge of AI, because they would simply prefer less regulation in their field, or for another reason. In any case, these findings suggest that there may be a contrast in opinion about the importance of regulation for AI between those creating these technologies, and the public, who will inevitably be impacted by their decisions.

Strong public opinions around privacy and governance are also undermined by the fact that even when privacy issues come to light, these companies remain extremely popular. Despite the fact that privacy is a primary concern in the digital age, these attitudes about privacy are sometimes inconsistent with actual behaviour, which is referred to as the 'privacy paradox' (Kokolakis, 2017, p. 122). Although there is extensive research around whether privacy concerns impact actual behaviour, results have been contradictory and it remains an unresolved issue (p. 132).

Although the privacy paradox is still a contentious topic, several studies have identified factors that are likely to impact people's willingness to adopt a technology, and share personal data. For example, several researchers have examined how perceived social benefits outweigh privacy concerns when it comes to the use of social media sites (Debatin et al., 2009; Lee et al., 2013). Additionally, some argue that the perceived and real life benefits of using the internet are generally more influential "factors on consumer's decision making than online privacy concerns" (Anic et al., 2019, p. 813). It is, however, not yet clear whether these kinds of perceived social, or real-life benefits are a strong motivating factor for people when it comes to embracing VA technologies specifically.

Apart from social benefits, other factors that can impact how much people are willing to disclose and share online, are enjoyment and trust. Positive mood-inducing features can have a significant effect on user trust and can lead to greater information disclosure (Wakefield, 2013). A study that involved interviewing VA users found that consumers had developed stable interpersonal trust relationships with a parental character they had attributed to the device, which impacted their trust and willingness to use these technologies (Foehr & Germelmann, 2020).

Perceived usefulness of VAs is another key factor in people's willingness to adopt these technologies. A recent study showed that those who believe these systems are useful cared less about privacy concerns, but for those who did not consider VAs useful, privacy concerns tended to have a significant impact on willingness to adopt the technology (Vimalkumar et al., 2021).

Recommendations and Guidelines

With big tech coming under growing scrutiny, there may be increased pressure for them to address ethical concerns. An example of a set of ethical guidelines made by the company itself is Google's AI principles, which state that Google beieves that AI should be socially beneficial, built for safety, be accountable, not reinforce bias, and several other guidelines (Google, n.d.). Microsoft AI principles are similar, and include: fairness; reliability and safety; privacy and security; inclusiveness, transparency, and accountability (Microsoft, n.d.). Despite the growing emphasis on ethical responsibility, it is hard to tell if the publication of ethical values or principles by tech companies is simply 'ethics-washing' where companies perform ethical responsibility, or if they are indeed engaging meaningfully with actual ethical issues (O'Reilly Media, 2020). Notably, the strength of these companies' monopolies and their control of the

market also diminishes consumer power to influence their practices and potentially makes them less motivated to follow through on meaningful ethical evaluations or interventions.

The concentration of data and power by large tech companies like Apple and Google, combined with lagging government response, has led many groups to develop recommendations for regulation. There is no shortage of principles or guidelines for the ethical development of AI. A 2020 study compared 36 prominent AI principles documents and identified eight key themes across these guidelines (Fjeld et al., 2020). These main themes were privacy, accountability, safety and security, transparency and explainability, fairness and non-discrimination, human control of technology, professional responsibility, and the promotion of human values (p. 13). Although these principles were found to exist in many of the documents, there remains no global standard for implementation of these ethical recommendations and there does not appear to be significant structures for accountability. It would also be difficult to ensure compliance to recommendations and hard to regulate what happens behind closed doors at large powerful technology companies. Notably, anthropomorphization and potential for manipulation were not major themes in many of the documents, demonstrating a need for further research and attention in these areas.

Although concern around manipulation by human-like AI is not central to many guidelines for ethical AI development and regulation, there are several notable recommendations that have been made. For Social AI like VAs, the IEEE (2019) suggests in *Ethically Aligned Design*, that "intimate systems must not be designed to explicitly engage in the psychological manipulation of the users of these systems unless the user is made aware they are being manipulated and consents to this behavior" (p. 96). Scheutz (2012) suggests that it should be law that AIs should "continuously signal, unmistakably and clearly, to the human that it is a machine, that it does not have emotions, that it cannot reciprocate" (p. 218). Value Sensitive Design (VSD) is another approach to the ethical design of technology that could be applied to VAs. VSD suggests that technology should be designed in a way "that accounts for human values in a principled and comprehensive manner throughout the design process" (Friedman et al. 2013, p. 55). That said, critics of the VSD approach argue that it lacks a methodology for distinguishing genuine moral values from mere stakeholders-preferences (Jacobs & Huldtgren, 2018).

Anticipated Contribution

A great deal of criticism and recommendations for AI have been produced by scholars, international organizations, companies, and other interest groups (Fjeld, J. et al., 2020). There is, however, very limited research that approaches issues related to ethics of VA technologies that incorporates public opinion data. Where public opinion data has been used in the research of VAs, it is often from a marketing perspective and related to consumer willingness to adopt these technologies. As discussed above, public opinion has also been used to research privacy concerns more broadly and phenomenons like the privacy paradox; however, the relationship between public opinion, anthropomorphization, and ethical concerns around VAs is currently underexplored in the literature.

As developers continue to improve Social AI technologies, the ability to impersonate human behaviour, personality, communication, and relationships may have significant social impacts. As Crawford (2021) points out, AI has been designed to reinforce existing structures of power and "countering that requires centering the interests of the communities most affected" (p.224). This research will therefore aim to build off of public opinion data to explore these issues and identify future research areas.

Research Questions

This research addresses the following questions:

1. What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?
2. What can be determined by evaluating public opinion about if and how Virtual Assistants should perform humanlike characteristics and behaviour?

Chapter 3: Methods

This exploratory research used non-probability sampling and an online survey to gather public opinions about Virtual Assistants (VAs). Results of the survey were summarized and analyzed using a descriptive statistics approach, in order to identify trends and provide recommendations for future research around VA technologies.

Development of Research Questions

Given the popularity, rapid development, and societal implications of VA technologies, there were many possible approaches that could have been taken to explore the ethics of VAs. The initial plan for this research was to explore if and how companies that develop Social AI like VAs, create ethics policies or protocols, and whether those policies aligned with public opinion. The intention with that approach was to address the lack of accountability that currently exists when it comes to integrating ethics into Social AI development; however, due to the fact that policy development information for these kinds of technologies was not readily available and that publicly available information would be impossible to verify without access to the inner-workings of these companies, the focus of the research shifted away from policy development.

There was not substantial literature in the area of public opinion around VAs, especially in regards to the performance of human-like characteristics from a non-marketing perspective. In light of this, an exploratory survey was created in order to capture opinions on different aspects of VA technologies including ethical concerns, trust in VA companies, VA use, and opinions regarding the portrayal of human-like characteristics by VAs.

Given that this research was exploratory both of the following research questions were made intentionally broad:

- What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?
- What can be determined by evaluating public opinion about if and how Virtual Assistants should perform humanlike characteristics and behaviour?

The survey was used to cast a wide net to gather opinions around ethical concerns, and human-like characteristics from multiple angles.

Survey Development

The survey used for this research (Appendix A) was an anonymous, voluntary, publicly accessible, online survey. The development of the survey was an iterative process that involved integrating feedback from the thesis supervisory committee, the delegated reviewer from the University of Alberta Research Ethics Board (application ID: Pro00110059), as well as graduate student peers and friends. The original survey was drafted after conducting the literature review and identifying the research questions.

Survey Structure

The survey was organized by grouping related questions together. The finalized survey consisted of the following four sections:

- The first section of the survey focussed on respondent demographics and gauging their familiarity with and use of VA technology. This section gathered basic demographic data about participants (age and gender). It also included questions about VA use, including how frequently respondents interact with VAs, which ones they have interacted with, where they use them, and the reason for VA use. Additionally, respondents were asked to rate their own awareness about how VAs collect and use data. This section also had participants describe using VA technologies and express what information they would rather get from a VA or another person. The questions included in this section lay the foundation for applying the results of the survey to the research questions.
- The second section directly aligned with the first research question and focussed on determining if participants had ethical concerns about VAs. This section gathered information on what concerns respondents have about VAs and to what degree they are concerned about various issues including transparency, VAs taking human jobs, manipulation, surveillance etc.. This section also asked participants to rate their overall level of concern around VA technologies and their overall level of concern about technology in general. Additionally, participants were asked if they trusted that technology companies who make VAs are ensuring that their technologies are not having a negative impact on people, society, and the environment.
- The third section focussed on the second research question and the portrayal of human-like characteristics by VAs. Participants were asked how human-like VAs should be, and whether VAs should have to identify themselves as AI. Respondents were also

asked for specific information about what gender, age, and personality traits VAs should perform.

- The last section collected opinions about data collection and who should be responsible for ensuring new technologies do not have negative impacts. This section included questions peripherally related to the two research questions that could provide interesting data to compare with or help contextualize responses in other parts of the survey.

In total, respondents were asked 61 questions. Almost all of the questions were mandatory for participants to answer, with the exception of a few short-answer questions that provided opportunities for participants to elaborate on their previous answers.

Gathering Feedback

This feedback stage of survey development was invaluable. Feedback was gathered from the supervisory committee and the researcher's peers. Data from the peers and friends who piloted the survey was not included in the data used for analysis, because changes were made to the survey after they had tested it. The feedback gathered during the piloting stage led to important changes in the wording and framing of questions and the inclusion of several questions that were not in earlier versions of the survey. Several examples of changes to the survey that came out of the feedback stage include:

- The addition of "It depends" as an option for several questions, followed by another question that allowed respondents to elaborate on their answer. This provided greater detail in responses.
- The addition of a "Not sure" option for several questions, including a question asking whether respondents use location services on their phone.
- The inclusion of two positive questions, 'What do you think is the most positive thing about Virtual Assistant technology?' and, 'What do you hope Virtual Assistants will be like in the future?' These questions were included to help balance the tone of the survey, which includes many questions about the concerning aspects of these technologies.

Deciding Who Could Participate

Very few limitations were put on who could respond to this survey. The only criteria was that people needed to be 18 or older and they needed to have a basic understanding of what VAs are and how they are used. Although those under 18 may have valuable opinions on the topic of

VAs, gathering data from that age group was not necessary to answer the research questions and the intention of this research was to focus on the opinions and attitudes of adults.

Respondent Anonymity

No information was collected which could be used to identify participants. It was made clear in the consent letter that preceded the survey, that respondents should not feel any pressure to participate. Because respondents could not be identified, confidentiality was automatic and the anonymization of data was not necessary. The choice to have the survey be anonymous was appropriate given this research was not focussing on one particular group of people, whose membership to a particular group needed to be confirmed. There was also no need to be able to contact or follow up with respondents for this research, although contact information for the researcher and supervisors was provided if respondents had questions.

Having the survey be anonymous also served to make the survey more straightforward from an ethics perspective. The survey was partly distributed using personal contacts and having responses be anonymous may have made those who know the researcher personally feel more comfortable sharing their opinions, and may have helped encourage candid responses.

One disadvantage of an anonymous survey is that, because no identifying information is recorded, it is not possible to remove a participant's responses from the data after they have completed the survey. Respondents were informed of this in the consent letter prior to completing the survey, and so far no respondents have reached out with questions or concerns. Another potential problem with anonymous surveys was that there was nothing that prevents someone from submitting more than one response. This scenario seemed unlikely, however, given that there was no direct incentive to completing the survey and therefore it is hard to imagine what would motivate someone to complete it multiple times. There was also no indication that anyone submitted responses more than once, and no two respondents gave identical answers.

Consent Letter Development

The consent letter for this research (in Appendix A) was developed by consulting the sample consent letter provided by the University of Alberta Research Ethics Board and past consent letters for different research projects. The fact that the survey was anonymous and voluntary was stressed in the consent letter. The consent letter also included a brief definition of what VAs are and examples of popular VAs (Siri and Alexa) to ensure that participants

understood what the research was about. The letter included a summary of the aim of the research and contact information for the researcher, the research supervisors, and the ethics board, in case participants wanted more information.

Selecting a Platform

Hosting and distributing the survey online was a secure, anonymous, and convenient way to gather opinions about VAs. This research looked at public opinion generally rather than collecting data from any particular group. It was therefore important that the survey was easily accessible online and could conveniently be shared. Several survey hosting sites were considered including Qualtrics and Survey Monkey, but in the end the decision was made to host the survey through GoogleForms. GoogleForms was selected due to its ease of use, reputation for being secure, affordability (free through University of Alberta subscription), ability to export responses to a spreadsheet document, and because access to the data was password protected. The fact that the responses would be hosted online also meant that data would not be lost if the researcher's computer crashed. Once ethics approval for the research had been attained, a small group of people, consisting primarily of the researcher's peers, were asked to test a copy of the survey to make sure that questions were clear and that responses were recorded successfully in GoogleForms before the survey was made public.

Survey Distribution

The survey was shared in Canada and particularly in the province of Alberta; however, those from elsewhere were not prohibited from participating. Once ethics approval was obtained participants were recruited through a combination of listservs, social media promotion, and by word-of-mouth or snowball sampling. As Vehovar et al. (2016) state, it is essential when using volunteer sampling to spread the non-probability sample as broadly as possible by combining various recruitment channels (p. 332). A distribution plan (Appendix C) was created and followed to circulate the survey. The distribution strategy included a variety of different social media pages and personal contacts that aided in reaching different audiences.

Another recommendation Vehovar et al. (2016) make is to shape non-probability samples so they "reflect, as much as possible, the structure of the survey population" (p. 333). To apply this recommendation, respondent demographic data was examined part way through the data collection process so that age and gender groups that were underrepresented in the data at that point could be sought out during a second round of distribution. In order to determine which ages

were under-represented, respondent ages were compared with age data from the 2016 Alberta census, not including people under 18 years old (Appendix B). In terms of gender representation, the aim was to have roughly equal data collected from cisgender males and females, and at least some representation from other genders.

When respondent demographics were analyzed halfway through data collection it became clear that men and older respondents were not responding as frequently as women and members of the younger age categories. Attempts were made to recruit these underrepresented groups by reaching out to contacts that were members of these groups and asking for their assistance sharing the survey. These efforts did result in higher percentages of men and older respondents, than there was initially; however, these groups were still underrepresented in the sample. This is discussed in greater detail in the following chapter.

Analysis

90 respondents completed the survey and submitted their responses. The analysis process for this research followed the descriptive statistics approach of describing, organizing, summarizing, and discussing trends in the data (Coleman, 2018, p. 488). The analysis process involved these procedures: 1) describing, organizing, and summarizing responses, 2) using cross-tabs to explore the relationships within the data, 3) conducting additional research to interpret the results, 4) considering possible explanations for responses, including thinking critically about the framing of questions, and, 5) identifying areas that require further research.

Describing and Organizing Responses

The first phase of analysis was to carefully review all of the responses to each question in the survey. GoogleForms, which hosted the survey, automatically produced charts and graphs of responses. These visualizations helped provide a preliminary understanding of the trends in the data and areas where there was strong consensus among respondents.

Following a preliminary review of the responses it was decided that results should be discussed over three chapters. The first analysis chapter would outline respondents' demographics and VA use, and the following two chapters would focus on applying responses to the research questions. Following this structure, the questions were organized into three categories: 1) those that provide insights into respondents and their VA use behaviours, 2) those that relate to ethical concerns around VAs, and 3) those that relate to the portrayal of human-like characteristics by VAs. After organizing the questions into three categories, the analysis and

interpretation process began by describing the results of all the questions in full sentences and making notes about possible relationships to consider and explore further. This stage also included the creation of tables, graphs, and pie charts to visually represent some of the responses.

There were several questions that did not naturally fit into any of these three chapters and were therefore not included in the discussion. Examples of questions that were not discussed in the analysis include a question about when it is appropriate for law enforcement to be able to access personal information gathered by a VA, and a question about how respondents want information about how your data is managed to be communicated to them. These questions were only peripherally related to the research and in the end were not necessary to pursue as part of the analysis.

Using Cross-tabulations

An important part of the analysis was investigating possible correlations in the survey responses. Cross-tabs were created using the pivot table function in GoogleSheets. In total 15 pivot tables were created to locate relationships in the data. One example was looking at whether a relationship existed between respondents' overall level of concern about VA technologies, and how often respondents interact with these technologies. This and other relationships, or surprising absense of relationships, are discussed in the analysis and in some cases presented in charts or tables to illustrate comparisons.

Although creating cross-tabs was very effective for examining certain relationships within the data, especially where there were clear variables to compare, there were many instances where large text filled answers could not be analyzed easily using this method. Other built-in functions of GoogleSheets were used to examine the relationships in the data in these cases. For example, in some cases responses were filtered so that only answers containing certain words appeared, the remaining responses were then compared with answers to other questions to find possible relationships. This technique was especially useful for questions where respondents could select multiple responses, such as the question of what personality traits participants believe that VAs should have.

Interpreting Responses with Additional Research

This interpretation process involved speculating on the potential factors that might have contributed to different responses, inferring what the responses might indicate, and conducting additional research to aid in the explanation of trends in the data. Where additional research that

could inform interpretations was not found, ideas for future research that would explore the topic further were suggested.

Additional research was also used to contextualize the analysis. This is particularly true for the discussion of the portrayal of human-like characteristics by VAs in Chapter 5. In researching the portrayal of personality and human traits such as age and gender, the idea of VAs as characters and the performance of servitude emerged as important themes. These ideas are explained at the beginning of Chapter 5 and are returned to throughout the discussion of survey responses. In this case previously conducted research is used to support the idea that voice carries a great deal of information and the character of a VA is designed with great intention.

Accounting for Research Limitations and Identifying Areas for Further Research

The limitations of the research are discussed throughout the analysis. For example, limitations related to the sampling method are discussed in the next chapter which describes the respondent demographics. Care was taken with the wording of the analysis to continuously signal to readers that these findings are not generalizable. In several instances the precise wording of the question is considered a potential factor that could have influenced responses. In most cases further research is necessary to pursue these interpretations in greater detail. The limitations of what can be determined through this research and how it may or may not relate to previous research, is clearly outlined throughout the discussion and suggestions for research that explore these ideas in greater detail are provided. A summary of areas for further research is discussed in the conclusion (Chapter 7).

Chapter 4: Respondents and Virtual Assistant Use

The following three chapters present the results of the survey as well as relevant analysis and interpretation. The findings and discussion will be organized as follows:

- Chapter Four, "Respondents and Virtual Assistant Use," provides a foundation for applying the results of the survey to the research questions. This chapter begins the analysis by discussing respondent demographics and the limitations of this research. The rest of the chapter aims to gauge the level of familiarity respondents have with these technologies as well as their general attitudes toward Virtual Assistants (VAs). This chapter will therefore also include a summary of if and how respondents have interacted with VAs and which VAs they are familiar with. In order to provide an overview of respondents' attitudes and information seeking preferences as they relate to VAs, questions about how participants describe using VA technology, and what information they would rather get from a VA than a real person, are also discussed in this chapter.
- Chapter Five, "Ethical Concerns About Virtual Assistants," applies the survey results to the first of the two research questions: *What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?* This chapter discusses the results of the survey questions related to trust, responsibility, privacy and transparency.
- Chapter Six, "Opinions on the Performance of Humanlike Characteristics," applies the survey results to the second research question: *What can be determined by evaluating public opinion about if and how VAs should perform humanlike characteristics and behaviour?* This chapter focuses on interpreting the respondents' opinions on the portrayal of characteristics such as age, gender, and personality traits by VAs. It also discusses how human-like participants think VAs should be, and whether a VA needs to identify itself as AI.

Demographic Information and Limitations of the Sample

The survey (Appendix A), accepted responses between July 15th, 2021 and August 26th, 2021, gathering a total of 90 submissions. The survey data was collected through a self-selection process and uses non-probability, volunteer sampling. As discussed in the methods chapter, this survey was circulated through personal contacts, social media, and word of mouth, and therefore results are not generalizable to any particular population. As suggested by Vehover et al. (2016), any trends that appear in the data will be clearly represented as 'indications' rather than

'estimations' of public opinion (p. 344). The findings can inform further research, discussion, and investigation.

At the time of data collection, 95.6% of respondents lived in Canada. Although the sample is primarily composed of those living in Canada, VA technologies have become popular around the world and opinions around these technologies likely vary considerably between regions and cultures. The data collected here should not be considered representative of global perspectives. Even though the vast majority of participants were based in Canada, results should also not be considered a snapshot of Canadian public opinion either. No racial, religious, political, or economic information was gathered from participants and so it is impossible to determine whether the sample accurately reflects the diversity of the Canadian population. Given that snowball sampling and word of mouth was used to conveniently gather this data it is very unlikely that respondents represent a cross-section of Canadian society.

Although no other demographic information was gathered, the survey did collect the ages and genders of participants. As will be discussed further in the following paragraphs, the data disproportionately represents the opinions of those between the ages 18-39 (66.6% of respondents). In terms of gender balance, men are slightly underrepresented, making up only 34.4% of the total respondents. Both of these factors should be considered as possible explanations for trends in the results. For example, the disproportionate representation of those under 40 may have contributed to trends in the responses, such as the large percentage of participants who have interacted with a VA (88.9%).

Age Distribution

As mentioned previously, despite efforts to circulate the survey amongst older generations, the age distribution of participants does not accurately reflect the population of Alberta (Figure 1) which was the original aim. The 2016 census data for the Alberta population is used here to help determine how the ages of the sample compare to the ages of a larger Canadian population (Appendix B). The province of Alberta was selected specifically because it is where many respondents are likely to inhabit given that the researcher is based in Alberta and the survey was distributed partially through personal contacts.

Of those who participated, 43.3 % were ages 30-39, 23.3 % were ages 18-29, 14.4% were ages 50-59, 10% were ages 40-49 , 7.8% were ages 60-69, and only 1.1% were 70 years of age or older. The consent letter that preceded the survey stated that respondents needed to have a

basic understanding of what VAs are and how they are used. It is possible that this requirement may have deterred older participants, who might be less familiar or confident with emerging technologies, from completing the survey.

Figure 1

Age distribution of participants compared to the 2016 population of Alberta (2016b).

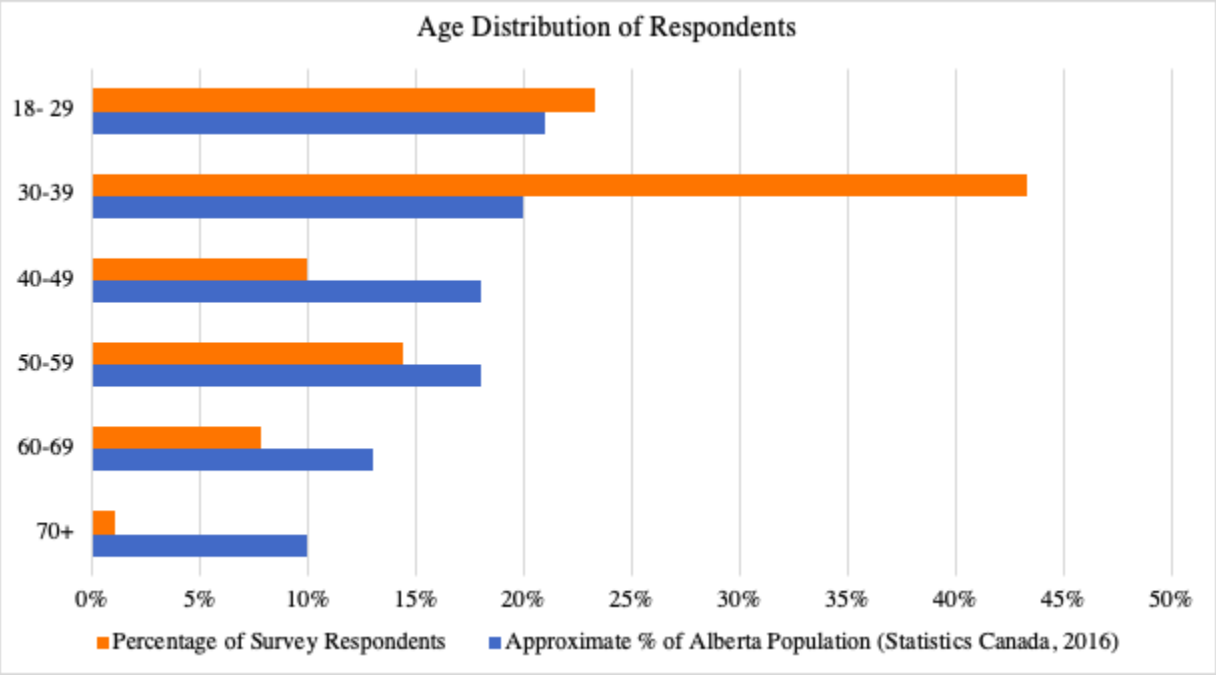

Gender Distribution

There were nine options for gender provided in the survey. The decision to include this many options for gender was made, in part, to provide inclusive options for respondents. The survey also included text stating that gender information would not be used to identify participants, which was a potential concern brought up during the ethics review. Gender information was also gathered in order to compare the gender of respondents with their preferred gender for VAs to portray, which is discussed in Chapter 5.

When asked to select their gender 56.7 % of respondents selected "Female," 34.4% selected "Male," 6.7% selected "Non-binary," 1.1 % selected "Transgender Woman," and 1.1%

opted not to provide a gender. The other gender options were not selected by any respondents. Unlike with age distrubution, the Alberta 2016 census data cannot be used to compare gender distribution in the responses, because the census collected sex information rather than gender information (Statistics Canada, 2016a). The aim when collecting responses was therefore to collect roughly equal numbers of responses from men and women as well as several responses from those who do not fall into either of these two categories. The survey successfully gathered opinions from a few non-cisgender people, but less success was found recruiting men than women. Other research has observed that women participate in survey research more often than men (Smith, 2008); however, the reason for this is unclear.

Virtual Assistant Use by Participants

Given that the survey was open to the general public, a brief assessment of participant familiarity and use of VAs is an important first step in analysis. This information helps contextualize other responses to the survey questions and how survey findings can be applied to the research questions in subsequent chapters. The following section summarizes if and how respondents have interacted with VA technologies, how frequently they use them, where they are used, and the reasons for using them.

Familiarity

The consent letter for the survey stated that participants "do not need to be a regular user of these technologies to participate" but, "should have a basic understanding of what they are and how they are used" (Appendix A). A simple definition of how the research defines 'Virtual Assistants' was also provided in the consent letter. The majority of people who completed the survey, 88.9%, have experience interacting with a VA. Presumably then, the other 11.1% of respondents have not personally interacted with a VA but have awareness of what they are.

The most popular VA was Apple's Siri, with 65.6 % of respondents saying they had interacted with Siri. 45.6% had interacted with Google Assistant, 27.8% with Amazon's Alexa, and 12.2% with Microsoft's Cortana. Samsung's Bixby was mentioned by 2 respondents using the "Other" option. "Ford Pass," "ATB V", and "Rogers VA" were mentioned once each. 43% of respondents have interacted with multiple VAs. A recent study of people from the US, UK, and Germany, suggested that Siri was more popular amongst those 18-24, whereas Alexa was more popular amongst older generations (Vixen Labs et al., 2021). This trend was not found among

respondents in this case however, with Siri being the most popular VA among respondents in every age group.

Overall the level of familiarity among respondents was very high. As mentioned previously, the survey responses came largely from those 18-39 years of age (66.6%) which could also explain the large percentage of people who have interacted with a VA. Another possible explanation is that regular users of Virtual Assistant technologies may have been more interested in the topic and motivated to complete the survey.

Frequency of Use

Although the majority of respondents had experience interacting with a VA at some point, the frequency of use for VAs varied considerably amongst respondents. There was a relatively even distribution between the number of respondents who have never used a VA (10%), those who used to use one but no longer do (13.3%), those who use one multiple times a day (12.2%), those use one once a day (11.1%), those use one several times a week (13.3%), and those use one roughly once a week (11.1%). "Roughly every two weeks" was selected by 6.7% of respondents and "once a month or less" by 18.9%. One participant had inconsistent responses, selecting both that they had "never used a VA" during an earlier question and that they used one "once a month or less" for this question. Because of this, the number of participants who selected that they have never used a VA for this question was 10% rather than 11.1%. Three participants used the "Other" option to elaborate on the circumstances in which they have used VAs, one stating that it was used only rarely and when convenient, and the other two stating that they used only to test it when they first got their device.

These findings suggest that the sample includes those who use VAs regularly, occasionally, rarely, and those who never use them. Although it was not done intentionally, the survey has managed to capture a range of familiarity levels with VA technologies despite being imbalanced in other ways such as age and gender distribution. 47.1 % of women, and 58.2% of men who responded use a VA once a week or more.

Location of Use

The most popular location for Virtual Assistant use was at home, which was selected by 67.8% of participants followed by "in the car" which was selected by 40%. As will be discussed in the following sections, some of the main benefits of VAs that participants identified are convenience and hands-free capabilities. This is undoubtedly a reason behind cars being a

popular location for VA use. Arguably hands-free capabilities are less of a factor for VA use being popular in the home, although they could come in handy when users are engaged in household tasks such as cooking, cleaning, childcare, etc..

Privacy concerns or feelings of self-consciousness may also inform where people prefer to use VA. Previous research has determined that people prefer to use VAs in private locations rather than public spaces (Easwara Moorthy & Vu, 2015). Other explanations for why the home could potentially be a popular space for VA use include the emergence of new smart home technologies and the fact that the home is simply where people spend a significant amount of time. These are only speculations however, as respondents were not asked what motivated them to use VAs in these spaces.

Reasons for Use

Although respondents were not asked for reasons why they use VAs where they do, they were asked "What do you use Virtual Assistants for?" Participants were prompted with 12 possible answers to this question and could select all that apply. There was also an "Other" option, which was used by 12 participants. The most popular activities respondents used VAs for were "Searching the internet" which was selected by 46.7% of participants and "Getting directions" which was selected by 41.1%.

"Fun (asking funny questions etc.)" was the third most popular option and was selected by 37.8% of respondents. Details were not gathered as to the exact nature of the fun, whether it was the unpredictability of the responses, the novelty of talking to a VA, or some other form of amusement. The question of what makes VAs fun might be significant. Future research could explore fun as a motivation for using VA technologies in more detail, including whether this enjoyment comes from the novelty of interacting with a VA or whether it contributes to a lasting positive user experience that motivates regular use of VAs.

A 2019 study found perceived enjoyment is more powerful than perceived usefulness for predicting people's intentions to use VA devices (Yang & Lee, 2019, p. 81). Further research could also explore the relationship between the enjoyment of VAs and other factors such as trust in these technologies.

Human or VA?: Information Behaviour Preferences

In order to get a better sense of respondent preferences when it comes to VAs and to capture an information behaviour component for this research, participants were asked what

information they would rather get from a VA rather than a real person. The findings indicate that there are several categories of information that the majority of respondents would rather get from a VA than from another human. The most popular answers for this question were "Directions" which was selected by 82.2% of respondents, "Simple facts (e.g. What is the capital of Manitoba?)" selected by 76.7%, and "Schedule details" selected by 52.2%.

Interestingly, only 6.7% of participants stated that they would rather not get any of the information options provided from a VA. There is clearly perceived value in VAs as information retrieval technologies, particularly when it comes to relatively straightforward queries. VAs provide a very convenient way to access the internet, which is where we often go for answers to simple questions or to fact check, rather than making an effort to consult with another person for easily accessible information. There may also be a feeling that this kind of information is something a person may remember incorrectly and that the internet might be a more reliable source for basic facts.

A limitation of the analysis is that the survey did not also ask participants whether they would rather get information from a VA than get it from a digital but non-voice activated technology, such as a search engine. It is therefore unclear exactly what motivates these preferences, whether it is the hands-free convenience and experience of using a VA in particular, or simply prefering to not interact with another human for these information needs. Human interactions are nuanced and complex, whereas VA interactions (at least for now) are more straightforward with no real social risk which may add to the appeal of this kind of interaction.

The survey also did not ask participants to explain their reasoning for preferring to receive information from either a human or a VA. It can be assumed that convenience could be a significant factor. Findings from other sections of the survey strongly suggest that convenience and hands-free capabilities are perceived as major advantages of VA technology. It requires very little effort to retrieve responses from a VA and if the information is not overly complex. As with the "Simple facts" and "Schedule details" options, there is a reasonable chance that the VA would retrieve accurate information. Getting directions from a VA would also be more convenient than pulling over to ask a stranger for directions.

In addition to convenience, there may also be social factors for preferring a VA. Further research could explore whether the social hesitancy we might experience when asking strangers for help might not be a factor when it comes to interacting with VA technologies. Perhaps using a

VA makes people feel less judged for not knowing a particular skill or could allow people to avoid feeling guilty for taking up a human's time.

Notably, this question was phrased; "What information would you rather get from a Virtual Assistant rather than a real person?" The survey question did not define exactly who that 'real person' is in this scenario. Responses may have been influenced by who respondents imagined this real person to be. Whether they were picturing this person as a friend, colleague, total stranger, professional, personal assistant, or someone else, likely impacted their answers. For example, it is possible that many people feel uncomfortable asking a stranger for information and would prefer the relative anonymity of asking a VA instead.

Interestingly, 43.3% of respondents stated that they would rather get "How-to information" from a VA rather than a real person. It is striking that so many respondents would prefer to get guidance on how to accomplish a task from a technology that, in almost every instance, would have no real-life experience performing that activity. Granted, the survey did not provide details on the level of complexity, or the importance of the task for which the user would be getting guidance.

Expertise, complexity, and risks of negative outcomes could be factors when it comes to whether or not to seek answers from a VA or another person. This is reflected in the responses related to medical information. The majority, 87.7%, of respondents would rather not get any medical or health related information from VA. Ten people (11.1%) responded that they would rather get basic medical or health related information from a VA, but only one participant stated that they would rather get complex medical or health related information from a VA. People are likely less comfortable getting medical information from VAs, because the risks of misinformation where health is concerned could be quite high. Medical information is also often very personal. Because of this there is an inherent affective aspect to medical consultation which is often expressed through the non-verbal cues we exchange when consulting with doctors.

The fact that there was hesitancy to ask for both basic and complex medical information from VAs, suggests a value in expertise as well. Generally speaking there is status and respect associated with doctors due to their education level. This high level of perceived expertise would be hard for any VA to compete with, despite the convenience and anonymity offered by a VA. A case study conducted in 2019 examining older adults' adoption of virtual doctors, found that a real doctor was prefered by older adults when they were seeking information specific to them

(Sin & Munteanu, 2019, p. 4). It was also noted that with human doctors there is a perceived ability to "identify a patient's real question, even if a patient lacked the knowledge to articulate it" (p. 4) which is a level of expertise not yet achieved by VA technologies in this context.

When it comes to asking for directions, however, perceived expertise appears to be associated with VAs, because the overwhelming majority of respondents would rather get directions from a VA than a real person. This is likely due to the integration of GPS technologies. Future research could explore if VAs programmed with specialized medical information retrieval, rather than those drawing from the internet more generally, would be perceived as having more expertise and be trusted to retrieve medical information to a greater extent than the popular VA options currently on the market.

Information behaviour research has also found that people do not want to spend a lot of effort seeking information for a task they will only complete once (Mawby et al., 2015, p. 477). The five most popular kinds of information that respondents would rather receive from a VA are 'Directions,' 'Simple facts,' 'Schedule details,' 'Translation services' and 'How-to information.' Many of the situations where users would seek information related to these categories would be one-off situations, such as getting directions to a restaurant or translating a word while traveling. In these situations people may prefer to use a VA because it is not worth the effort of seeking out from a real person when there will not be a continued use for that information.

How Participants Describe Virtual Assistant Technology

The next section of the survey aimed to capture respondent attitudes about using VAs. Participants were asked to respond to the question "Using Virtual Assistant technology is:" by selecting from 16 descriptive words. An 'Other' option was also included and there was no limit to the number of descriptors that respondents could select. Capturing general sentiment was a necessary step in order to get a picture of who was completing the survey. For example, if a huge number of respondents only selected negative descriptive words, that bias in the sample would have to be accounted for when applying the survey results to the research questions in subsequent chapters.

An almost equal number of positive and negative descriptors were provided as options (Table 1).

Table 1

How Participants Describe Using VA Technology

Positive (or neutral*) descriptors	**Negative descriptors**
Convenient (57.8%)	Creepy (40%)
Helpful (44.4%)	Annoying (28.9%)
Amusing (33.3%)	Frustrating (26.7%)
Easy (26.7%)	Difficult (7.8%)
Fine (no strong feelings)* (24.4%)	Draining (4.4%)
Fun (12.2%)	Boring (3.3%)
Satisfying (4.4%)	Confusing (3.3%)
Empowering (3.3%)	Nerve-wracking (2.2%)

The seven positive options included "Amusing," "Empowering," and "Helpful." Eight negative options included "Confusing," "Annoying," and "Creepy." One neutral option, "Fine (no strong feelings about it)" was also provided. Positive descriptors were selected 164 times by participants whereas the negative descriptors were selected 105 times. Of the respondents, 35.6% selected only positive or neutral descriptors for using VA technologies, 14.4% only selected only negative descriptors, and 48.9% selected a mixture of positive/neutral descriptors and negative descriptors. The responses to this question suggest that there was not a disproportionately negative or positive attitude toward the experience of using Virtual Assistants amongst respondents. Almost half of participants have mixed feelings about VAs and there is good representation from those with only negative and only positive or neutral things to say about using VAs.

The Creepiness of VAs

Some respondents' mixed feelings about VA use seems to stem from frustration that the technology is not working as effectively as they would like, whereas other negative descriptors, such as "Creepy," arguably come from a place that is more instinctual. The fact that the "Creepy" descriptor was selected by 40% of respondents suggests that feelings of discomfort, distrust, or unease are common when using VAs.

Tene and Polonetsky (2014) thoroughly examined the relationship between creepiness and technology and suggest that the word "creep" is commonly used to describe situations where our social values do not align with a new technology's capabilities (p. 60). They identify these five business models or activities that are labeled as creepy: ambient technologies, such as those where users' physical location is known; social listening; personalized analytics; data-driven marketing; and, new product launches, particularly if they involve unexpected uses of existing data (pp. 61-71). All five of these activities are within the capabilities of VAs. Additionally, all of these "creepy" activities can be linked to the idea that a VAs "knows" information that was not explicitly provided by the user. It is obviously outside of our social norms for another person to have information about us that we did not give them, such as knowing our location at all times, which is likely why we experience unease realizing that a VA has access to this data.

41% of female participants, 100% of non-binary participants, and 26% of male participants described using a VA as creepy. During their development of the Creepiness of Situation Scale (CRoSS), Langer and König (2018) hypothesized that women perceive more creepiness than men, and found this to be the case in a study where women were approached on the street (p. 11). The results of this survey appear to align with the theory that creepiness could be experienced more by non-cisgendered men in the case of VA use as well.

The Value of Convenience

The most popular descriptive word selected to describe using VA technology was "Convenient," which 57.8% of respondents selected. In a later section of the survey, when asked in a short-answer question what the most positive thing about VA technologies was, 20 respondents also mentioned the word "convenience" specifically. Relatedly, 17 responses to that same question mentioned hands-free abilities, ease of use, and saving time as being the most positive thing about VA technologies, which are all arguably related to convenience as well.

The convenience of VA technologies is undeniable. A VA can be directed to complete a task while the user's hands or eyes are preoccupied. They can also be interacted with in almost

any relatively quiet environment. The fact that these technologies can be communicated with verbally, the way that humans have evolved to interact with each other, also adds to convenience. For most people talking is instinctive and thus speaking with a VA does not require extensive technical know-how or even the ability to type well.

The convenience of VAs stems from the voice-activation and response capability. The listening and responding functions are what sets VAs apart from other information retrieval technologies and are essentially what gives these technologies value. As will be discussed in greater detail in the next chapter, these capabilities are also concerning, especially when users are uncertain about how these technologies work and the data collection practices.

Chapter 5: Ethical concerns about Virtual Assistants

The following chapter applies the survey findings to the first research question: *What ethical concerns, if any, do people have about Virtual Assistants and their data collection practices?* This chapter will discuss survey responses to questions related to ethical concerns, including those around trust in technology companies, data collection practices, and who should be responsible for making sure that technologies do not negatively affect people, society or the environment.

Concerns About VA Technology

In order to identify what ethical concerns participants might have regarding VAs, they were asked to rate their level of concern as "Not at all concerned," "A little concerned," "Moderately concerned," "Very concerned," or "Extremely concerned" for 20 common ethical concerns related to technology. These concerns were sourced, in part, from various literature identifying key ethical issues in AI development (Fjeld et al., 2020; IEEE, 2019). To identify which issues triggered significant levels of concern for participants, the percentage of respondents who were "Very concerned" or "Extremely concerned" for each issue was calculated (Appendix D, Table D1). Using this approach, the top ten issues that most commonly elicited significant levels of concern were identified (Table 2).

Table 2

The Ten VA Concerns Respondents Were Most Frequently "Very" or "Extremely" Concerned About

Specific Concerns	Percentage of respondents either "Very concerned" or "Extremely Concerned"
Lack of Transparency by Companies	67.8%
Manipulation by Companies	63.3%
Increased Political Polarization	60%

Corporate surveillance	55.6%
Political Manipulation	55.6%
Lack of anonymity (having your data identifiable to you)	54.4%
Risk of Privacy Breach	46.6%
Consumerism/buying things you don't need	41.1%
Inequality in access and usability of these technologies	41.1%
Risk of VA being hacked	37.8%

Key Concerns: Manipulation and Lack of Transparency

The two issues that most frequently elicited high levels of concern from respondents were "Lack of transparency by companies" and "Manipulation by companies" with 67.8% and 63.3% of respondents being very or extremely concerned about these issues, respectively. Arguably these two concerns go hand-in-hand. There is no such thing as transparent manipulation. "Manipulation" is defined as involving artful and unfair means to serve one's own purpose (Merriam-Webster, n.d.a). The potential for manipulation is far greater when coupled with a lack of transparency.

As mentioned in the literature review for this research, Shoshona Zuboff argues that the potential for manipulation marks a shift from traditional capitalism (Zuboff, 2015, p.75). It is understandable that manipulation by companies would be a major concern, especially compared to more obvious, traditional forms of advertising manipulation. Newer algorithm-based targeting methods, which are a fixture of the surveillance capitalism business model, are much more mysterious. The average person does not understand how the VA algorithm has been trained, what biases it might have, what data it is collecting, which third-parties have access to their information, and the degree that personal information is influencing the responses or recommendations generated by the VA. On the other hand, when someone watches a commercial

for a product or service, they know the company is trying to manipulate them. Even though commercials are aired on specific channels and times of day to reach target demographics, the advertisers likely do not have access to a lot of personal data about the specific individuals they are targeting.

Interestingly, unlike the high levels of concern expressed for manipulation by companies, only 28.8% of respondents expressed heightened levels of concern about threats to autonomy. Manipulation and threats to autonomy are two concepts that seem logically related, with manipulation being a key factor that could impact autonomy. This contradiction (Figure 2) may stem from a phenomenon observed by Cho et al. (2010), suggesting that people can display a strong optimism bias about online privacy risks, "judging themselves to be significantly less vulnerable than others" (p. 994).

This may also have been unwittingly supported by the wording of the question; this concern was worded, "Threats to autonomy (the ability to make your own decisions)." The section in parenthesis was added to provide a definition for "autonomy" for participants who might be unfamiliar with the term; however, the inclusion of the word "your" (although meant as a plural "your") may have made the response to this question more personal. Thus, respondents may have been impacted by an optimism bias when responding, considering themselves less vulnerable to manipulation than others, and thus, threats to their autonomy significantly less concerning than manipulation more generally.

Figure 2

Levels of concern for "Manipulation by companies" and "Threats to autonomy (the ability to make your own decisions)"

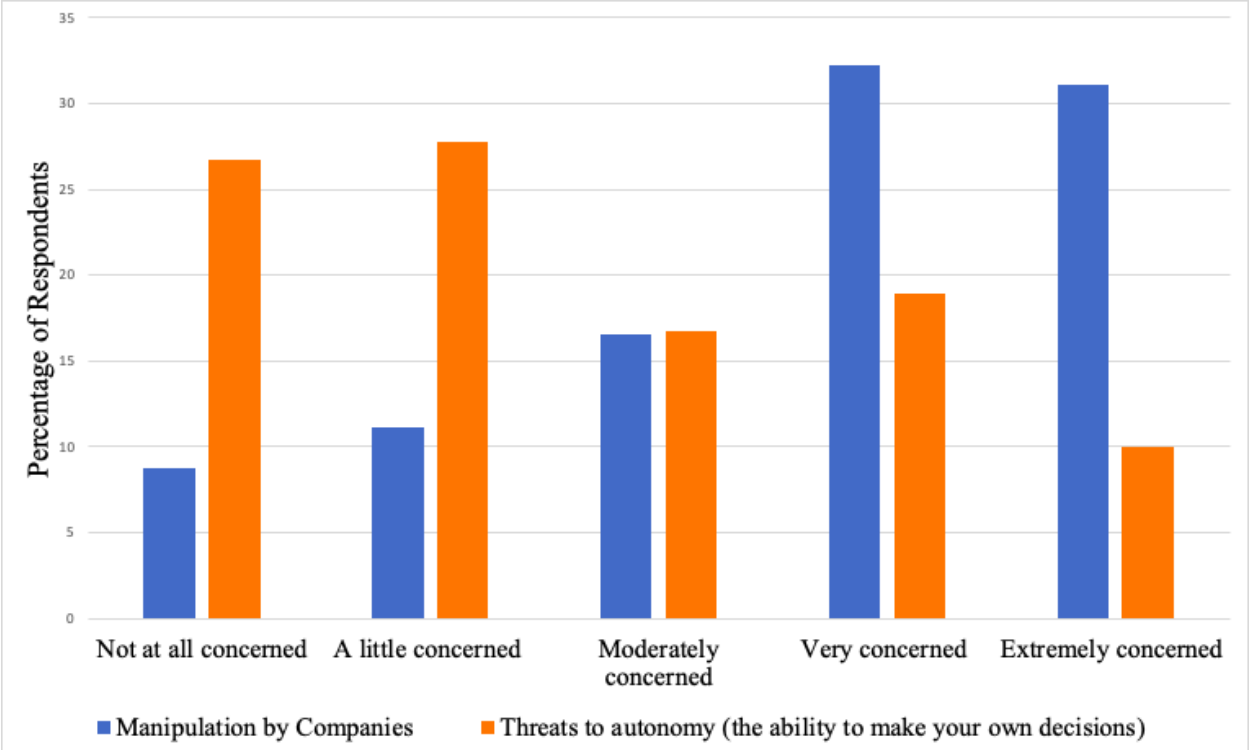

Key Concern: Virtual Assistants and Political Polarization

The third highest rated concern was increased political polarization. On the surface this seems like an issue that would be more directly associated with social media technologies rather than VAs, due to the potential impact an echo chamber effect may have on increasing polarization on social media platforms (Williams et al., 2015). Despite polarization being more directly associated with social networking technologies, there is growing interest in the role that VAs might play in political polarization as well. In a 2019 study, researcher Christopher Ojeda interviewed both tech workers, and VAs themselves, to assess how VAs are programmed to respond to political queries (Ojeda, 2019). Ojeda found that VA responses to questions related to political attitudes, appear to be pre-programmed by technology companies "rather than the result of an unimpeded machine learning algorithm" (p. 885). This aligns with other research that computer scientists and related technical experts often make the technical choices that have ethical and social consequences (Stahl et al., 2016; Moss & Melcalf, 2019).

Ojeda also found that although the economic considerations of companies may currently have more influence on the political leanings of VAs, in the future, responses "will likely create political echo chambers in which responses are highly customized to individual users" (p. 884).

That said, even when it comes to social media technologies, there are researchers who argue that the echo chamber phenomenon does not have as great an impact on polarization as is generally assumed (Dubois & Blank, 2018; Nguyen & Vu, 2019). Research has also found that whether social media leads to echo chamber-related polarization, depends greatly on the platform (Kitchens et al., 2020). Thus, it is unclear if VA technologies could contribute to political polarization in a significant way, even if they do eventually cause an echo chamber effect.

In a later question in the survey, participants were asked if the data collected by a VAs should be used to help curate people's social media feed. An overwhelming majority, 90%, of respondents said "No." People appear to feel very strongly that this is not an acceptable use of data. Although this question did not mention politics or political content specifically, the heated political landscape of the last few years and heightened concern about political polarization was likely a factor that contributed to the strong sentiment that our data should not be used to curate what we see on social media. This may also be reflected in how respondents rated their concern about "Political manipulation" which was also a significant concern for respondents.

Key Concern: Corporate Surveillance

Another issue that elicited higher levels of concern for respondents is corporate surveillance. Only 5.5% of respondents were not at all concerned about corporate surveillance. It is impossible to think of an example from the pre-internet era, of a private company having access to anywhere near the amount of personal information that modern tech companies collect. The significant concern for this issue may be infludenced by the unprecedented nature of this kind of mass data collection, and how quickly the practice has become commonplace. The level of concern expressed in the responses for "Corporate surveillance" was generally much higher than for "Government surveillance." However, this may be due to the fact that VAs are owned and developed by corporations and thus participants have less reason to be concerned that the government would be surveilling them with this particular technology.

Having participants identify their level of concern around corporate surveillance by VAs, was not the only question in the survey related to surveillance and data collection. When asked to rate their own awareness about how VAs store and use personal data on a scale of 1-10 (1 being "Completely unaware" and 10 being "Extremely aware") there were significant responses across the spectrum. 27.8% ranked their awareness between 1 and 3 on the scale, 26.7 % between 8 and 10, and the rest, 45.6%, placed themselves in the center of the spectrum by selecting a ranking

between 4-7 out of 10. The average rating was 5.4/10. The self-identified awareness level did change slightly with age, with an average rating of 4.2/10 amongst participants 50+, and an average rating of 6/10 for those 18-39. It may be that older respondents are slightly less confident in their awareness level of newer technologies, although it should be reiterated that only 21 responses were collected from those 50 years old or older. The results of this question indicate that there is a significant range when it comes to the degree of perceived awareness of data collection practices of big tech companies.

The existence of multiple questions related to the topic of data collection in the survey, provided an opportunity for results to be examined closer. A cross-tab was created to see if there was any relationship between respondents' self-identified awareness of corporate data collection practices, and level of concern around corporate surveillance. No strong correlation between these two variables was found. As an example, 60% of those who described themselves as completely unaware of how VA use personal data reported high levels of concern about corporate surveillance, compared to 66.7% of those who described themselves as extremely aware, and 62.5% of those who rated themselves in the middle of the scale. All of these figures are relatively close.

The fact that there does not appear to be a strong relationship between how aware a respondent believes they are about data collection and their level of concern for corporate surveillance is slightly surprising. This seems to indicate that corporate surveillance by VAs is a significant concern regardless of whether respondents feel they are aware of data collection practices or not. Future research could explore the relationship between awareness and concern in greater detail, perhaps using a larger sample size.

Of the 50 respondents who reported high levels of concern for corporate surveillance, 60% use location services on their phone. Although location is only one aspect of our behaviour that could be traced by corporate entities, and it is not known to what degree having their location tracked is particularly concerning for participants, it is interesting that so many respondents who are concerned about corporate surveillance use this tool. This comes back to the idea of the privacy paradox, which suggests that attitudes about privacy are sometimes inconsistent with actual behaviour (Kokolakis, 2017, p. 122).

That said, having your location tracked by your device seems to be a more obvious form of behavioural surveillance, which could make people feel more comfortable about it. Choosing

to keep location services activated on your device is a choice (except perhaps for the one respondent who was unsure about whether they used location services or not). Having the option to turn location services off at any time might ease some of the concern expressed by respondents.

Less Concerning Issues

The aim of this research, as outlined in the first research question, is to identify what concerns people have about VAs. The issues that respondents found less concerning, although interesting, will therefore only be discussed briefly. The three issues that elicited high levels of concern from respondents the least frequently were "VAs taking human jobs", "Technology dependency", and "Threats to autonomy (the ability to make your own decisions)" which has already been discussed in previous sections.

Of those surveyed, 57.8% felt no concern or "a little concern" about VAs taking human jobs. This may be because people tend to be more concerned about AI technology in the form of robots taking human jobs, rather than natural language AI replacing humans. Previous research on AI replacing human jobs has been dominated by the ways robotics and software can replace routine, highly structured, and repetitive tasks (Brynjolfsson & Mitchell, 2017). That said, the ways that AI development will affect employment in the future remain unclear and there may be significant developments that impact highly paid and educated workers as well (Muro et al., 2019). Although people may generally assume that AI will struggle to replace the social and interpersonal skills required for many jobs, it is hard to predict the ways that conversation technology and machine learning will impact the workforce. This may have contributed to the relatively lukewarm levels of concern around this issue amongst respondents.

Technology dependency was another issue that most respondents were not overly concerned about. It might be possible that an optimism bias, like the one observed by Cho et al. (2010) related to online privacy risks discussed earlier, exists in the area of technology dependency as well. Future research could examine this by comparing people's own perceived dependence on technology, with how they would describe the dependency on technology of those around them. It would also be interesting to research whether reframing the question to mention children, a group that users may feel would be more at risk for becoming overly reliant on technology, would influence responses.

Summary of Ethical Concern Responses

The results described above suggest that there are multiple issues with VA technologies that produce significant levels of concern for respondents. Concerns related to manipulation, corporations, transparency, surveillance, and politics were the issues where participants most frequently expressed high levels of concern. Issues related to autonomy and dependence less frequently elicited intense concern, perhaps suggesting that people are more concerned about corporations using mass-surveillance or large-scale manipulation, than they are about issues related to personal decision-making or behaviour. It could be possible that participants trust their own ability to perceive manipulation and counter a reliance on technology, but believe that these are concerning issues for society as a whole. Further research could use a larger sample size and explicitly frame concerns as having either personal or societal impacts to explore this in greater detail.

Trust in VA Companies

As discussed earlier in this chapter, lack of transparency by companies was the issue that most frequently had participants very or extremely concerned. Arguably, a lack of transparency would be much less concerning, if participants trusted that VA developers were making sure their technologies did not have a negative impact; however, most respondents do not trust that this is the case. Participants were asked if they trusted that tech companies who make VAs are ensuring that their technologies are not having a negative impact on people, society and the environment in three separate questions. For each question respondents were given the options of "Yes," "No," or "Somewhat."

For all three questions the vast majority of respondents selected "No" (Figure 3). 64.4% of respondents answered "No" to all three questions. This suggests that the majority of people feel companies are doing a poor job preventing negative impacts across the board. For the other 35.6%, the level of trust that companies are making sure their products do not have a negative impact, varied depending on if the question was focussed on people, society, or the environment.

Figure 3

Respondent Trust That Companies Ensure Technologies Do Not Have Negative Impacts

Survey Question	Responses
Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact people?	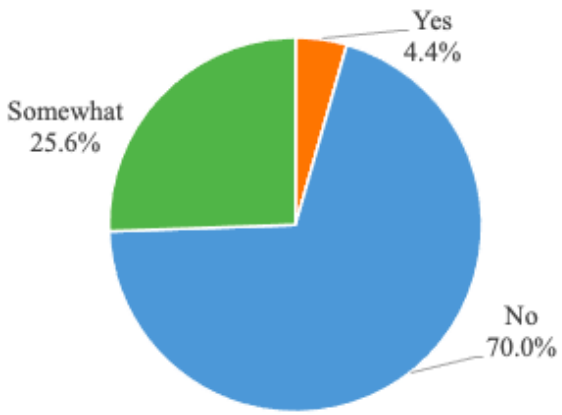
Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact society?	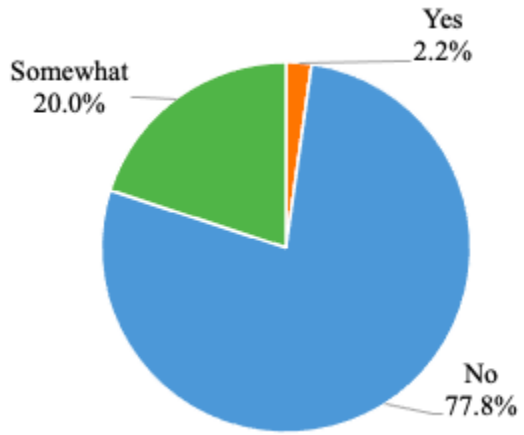
Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact the environment?	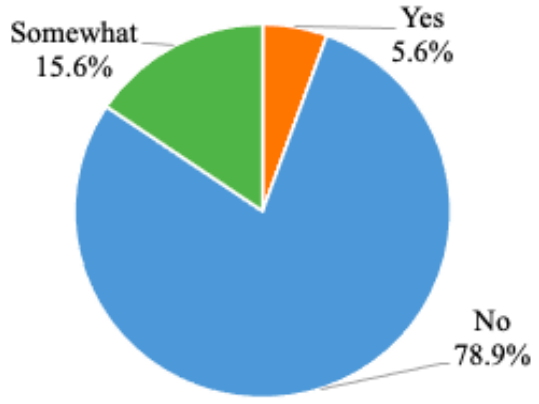

Ethical issues with AI, such as those related to privacy, accountability, transparency, and the environment, are issues discussed both in academic literature and in mainstream media (Morley et al., 2021, p. 240). The media coverage in recent years of various issues associated with AI has no doubt contributed to the lack of trust expressed by respondents by bringing greater awareness to these issues. Despite increased awareness and substantial research, there is still a significant gap between theory and practice within the AI ethics field (Morley et al., 2021, p. 240). Although there is no shortage of guidelines outlining principles for ethical AI development, (one report has so far identified 173 different AI ethics documents (Algorithm Watch, 2021)), applying these recommendations to design decisions is difficult. One major issue is that recommendations are either too vague or too restrictive to account for context, and for these reasons they are challenging to apply in practice (Morley et al., 2021).

In order to respond to this disconnect between theory and practice, some scholars recommend shifting the approach to AI ethics away from principles and toward procedural regularity and design strategies to make ethics more relatable to AI developers (Morley et al., 2021, p. 252). Perhaps better tools that translate ethics into practice can lead to ethical considerations being better integrated by companies and not used simply as a box to check for ethics-washing purposes. It is impossible to predict if and how these strategies will be adopted by tech companies and whether this would have a significant impact on public trust.

The matter of accountability, standards and regulation for such a quick moving industry is another huge challenge. When asked who should be responsible for making sure technology like VAs do not negatively impact people or society, only one respondent selected "Nobody." Equal numbers of respondents, 45.6%, selected "Governments" and "The companies themselves" (Figure 4).

Figure 4

Opinions on Who Should be Responsible for Preventing Negative Impacts of VAs

Notably, 87.8% of respondents selected more than one of the options provided, suggesting that the majority of participants feel this responsibility falls to multiple entities. Relatively few, 15%, of respondents felt that individuals should be responsible for making sure technology like VAs do not have negative impacts, but all of those who selected "Individuals" also selected other entities, meaning that they feel individuals should share this responsibility. These results suggest there may be public support for a layered or collaborative approach to preventing technological advances from causing harm. Luciano Floridi (2016) points out that the long-term impacts of technology can easily become everybody's problem and nobody's responsibility (p. 11). Floridi proposes that shared responsibility may have to become the "faultless responsibility," in order to shift the focus to the welfare of the system and away from the priorities and interests of the individual agents involved (p. 11). Morely et al. (2021) suggest that AI ethics would benefit from a customized approach, like those frequently found in research ethics and medical ethics, and involve a combination of "law, ethical governance policies, practices, and procedures, with contextual discourse and procedural support" (p. 252). They argue applying a similar combination approach to that used in other branches of applied ethics, could allow for good balance between regulation being too strict and too flexible.

Comparing VA Concern with Technology Generally

The majority of ethics-related questions in the survey focussed on VA technologies. In order to understand how the level of concern around VAs compared with that of technology in general, participants were asked to rate their overall level of concern around surveillance and data collection by tech companies in general, and by VA technologies specifically. 72.2% of

respondents reported the exact same level of concern for each. Of those whose level of concern was different for VA technologies, all but one of the responses shifted only one rating in either direction (e.g. from "A little concerned" to "Not at all concerned", or from "Very concerned" to "Extremely concerned"). These findings suggest that the level of concern experienced by respondents about surveillance and data collection is very similar for VAs as it is for technology in general.

Although respondents have not expressed substantially different levels of concern for the surveillance and data collection practices of VAs, there are striking trends in their opinions about some of VA technology's more unique capabilities. The next chapter will focus on the features that set VAs apart: their conversation capabilities and the portrayal of human characteristics.

Chapter 6: Opinions on the Performance of Humanlike Characteristics

The following chapter applies the survey results to the second research question: *What can be determined by evaluating public opinion about if and how Virtual Assistants should perform humanlike characteristics and behaviour?* This chapter focuses on interpreting respondent opinions on the portrayal of characteristics such as age, gender, and personality by VAs.

VAs can be considered anthropomorphic AI technology because, by communicating verbally, the performance of human characteristics by VAs is unavoidable. We naturally project features like gender, age, personality, education level, status, attitude, etc., onto these systems because we naturally assign these qualities, consciously or subconsciously, to other people. Although certain human-like qualities such as facial expressions and physical appearance are not portrayed by most VAs, a human-like voice carries a great deal of information. This information can be gleaned from vocal tone, pitch, dialect, speed of speech, vocabulary, inflection, emphasis, degree of articulation, and numerous other vocal features that influence the way we interact with each other. In the case of VAs, all of these vocal elements are designed with great care to elicit the desired reaction from users.

There are likely advantages to the absence of a physical depiction of a person in VA design. For example, the uncanny valley effect, where people are be put-off by an unnervingly not-quite-human appearance (Mori, 1970; Whitby, 2012), can be partially avoided by only using a human-like voice. With VAs, there is only one human-like communication method that designers need to get right, the voice.

Before proceeding with analysis, it is important to clarify that here, the use of the term 'human-like' is used to describe how a VA communicates and is perceived, not the actual intelligence mechanisms or procedures of the AI program. Russell and Norvig (2010), consider human-like AI as AI that was designed to replicate human ways of thinking (with all its nuance and flaws), and use the term "rational AI" to describe the much more common approach to AI design, which aims at an idealized version of intelligence where the machine is programmed to provide only correct responses (pp. 1-4). For the purposes of this research, however, "human-like" will not be used to describe the mechanisms of 'thinking' but rather, how responses generated by VAs (a "rational" AI technology) are presented in a way that mimics a human response.

The Degree of Humanity

As mentioned in the literature review, VA technologies, such as Google's LaMDA (Collins & Ghahramani, 2021), and Apple's Siri (Apple, n.d.) are rapidly improving their ability to impersonate human conversation. Most researchers in relevant areas unquestionably consider making these systems more human-like a desirable goal (Whitby, 2012, p. 244). Researchers and tech workers are not, however, the only groups of people that will be impacted by increasingly human-like AI becoming integrated into our devices and lives. For this reason the survey included the question, "How human-like should VAs be?" to see how people, who are not necessarily AI experts, researchers, or designers, feel about this (Figure 5).

Figure 5

Responses to the Question, "How human-like should VAs be?"

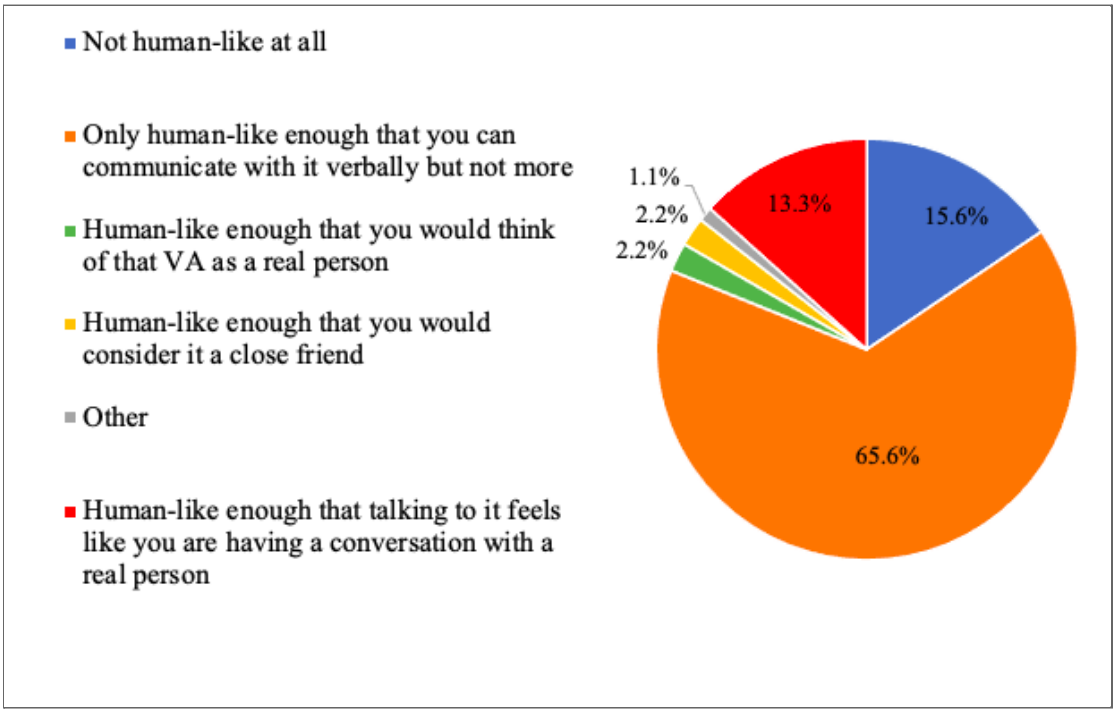

Contrary to the direction that VA technology is headed, respondents overwhelmingly believed that VAs should be minimally human-like. 65.6% of respondents believed that VAs should only be human-like enough that they could communicate with it verbally, but not more. 15.6% believe that VAs should not be human-like at all. Only 13.3% wanted VAs to be human-like enough that talking to it feels like you are having a conversation with a real person. It

is apparent in these results that public opinion, at least in this sample, does not align with the current trajectory for VA development. There were respondents who prefered VAs to have a great deal of humanity, however, with 2.2% respondents stating that they wanted to think of a VA as a real person, and another 2.2% wanting VAs to be humanlike enough that they could consider it a close friend.

It is possible that if this was the only question in the survey, responses would be different. Prior to being asked this question, respondents had answered 39 questions about VAs, many of which involved evaluating their level of trust in VA companies and rating their level of concern about these technologies. Perhaps by this point in the survey, participants were thinking more critically about the portrayal of human-like characteristics than they would have otherwise. That said, these results should not be disregarded. Arguably collecting the opinions of people who have taken a moment to stop and think about technology critically, provides important insights.

If indeed, when people are thinking critically about technology, the vast majority of those people's views strongly oppose the trajectory of technological development by the most powerful companies in the world, that disparity is significant. Further research must be conducted to compare public opinion to technical development trajectories in greater detail. A major challenge would be to encourage critical thinking among participants, while not embedding biases against VA technologies into the framing of questions.

In order to see if there was any relationship between respondent opinions on how human-like VAs should be and adoption of these technologies, a cross-tab was used to compare responses to this question with the frequency of respondent's interaction with VAs. Of those who believe that VAs should only be human-like enough that you can communicate with it verbally but not more, 31.1% use VAs at least once a day and 70% use them at least once a month. This suggests that even regular users of VAs believe VA should be minimally human-like.

VA as Character

The rest of this chapter will focus on the specific human attributes respondents believe VAs should perform. VAs are characters. These VA characters have been written slightly differently by the playwrights at Apple, Google, Amazon, etc.; however, just like in any good play, they serve a purpose in the narrative the company is trying to create. Brand anthropomorphism, which is defined as the perception of a brand as an entity with human-like features, and mental and emotional states, is recognised as an important construct in marketing

(Golossenko et al., 2020). VAs companies take the idea of brand anthropomorphism to the next level by putting great care into ensuring that VAs come across as a likeable character with good intentions that users can trust.

VAs as Servants

There are limited tools at a programmer's disposal when it comes to making sure VAs are perceived favourably by users. Without visual cues to make technology approachable it all comes down to the voice and what elements can be integrated into that voice that will trigger positive, familiar, reliable and trustworthy connotations and archetypes for users.

For companies, integrating humanness is a key factor in making VAs appealing and putting users at ease but only when combined with cultivating feelings of control, as previous research suggests that anthropomorphism decreases perceptions of risk for those who feel high levels of power (Kim & McGill, 2011). When people feel that their interaction may be risky in some way, human-like characteristics can actually enhance those feelings of vulnerability, but if they feel socially powerful, then anthropomorphization eases their perceptions of risk (p. 104). Applying this to VAs, perhaps the performance of servitude eases perceptions of risk when combined with the incorporation of human-like characteristics.

If users feel that VAs are there to assist them, and to respond appropriately to every demand, they might feel socially powerful and experience lower levels of vulnerability or concern around issues like surveillance or manipulation. This aligns with Shoshana Zuboff's (2019) observation that VAs allow people to feel like they have someone at their beck and call, a situation that creates the "sensation of mastery" when in fact we are "giving it away" (Zuboff, 2019, p. 260). Therefore promoting the idea of a VA is a servant or an assistant works to the advantage of companies trying to evoke a feeling of power and control in their users.

Scholar Justine Cassell suggests that Apple's choice to launch Siri with a male voice in Great Britain was made because the British have always had male servants and this was the company attempting to "evoke that stereotype of the always helpful, always present valet" (Sydell, 2018). This stereotype was used in a more obvious way by the question-answering service, Ask Jeeves (later changed to ask.com),which was named after P.G. Wodehouse's reliable and competent valet character (Sherman, 2003).

The Performance of Gender

Female by Default

With the exception of Siri's UK launch, when it first came time to cast the Virtual Assistant character, VA companies assigned the role to 'women.' Even if the companies do not admit it (Gartenburg, 2021), Siri, Cortana, and Alexa are female-sounding names with female-sounding voices. The reinforcement of gender stereotypes by assigning a female gender to a technology performing an assistant role has been widely criticized (King, 2017; Woods, 2019; Gartenburg, 2021; Sydell, 2018). The feminized, unthreatening personas that are often the default settings for VAs, are meant to establish trust between the technologies and the users. Siri and Alexa, for example, perform a feminine persona to enact "digital domesticity" that makes surveillance more "palatable for the mass, consumer public" (Woods, 2018, p. 344). Not only is gender used to make technology seem less threatening, research also suggests that female voice-based AI are perceived as more human than their male counterparts (Borau et al., 2021).

Although the first versions of popular VAs were all female-sounding, male-sounding voices are now usually provided as options. Even Amazon recently announced a male-sounding voice after holding out for a long time (Gartenburg, 2021). Although the female-sounding voice has almost always the default option, according to Apple (2021a) as of iOS 14.5 version of Siri no longer has a default voice and users are prompted to choose from the voice options provided right away.

Opinions on VA Gender

To compare the common practice of using feminine personas for VA technologies with public opinion, respondents were asked: "What gender should Virtual Assistants portray?" The survey options for this question were slightly different from the options for the question promting respondents to state their own gender. For example, "Transgender man" and "Transgender woman" were not given as options for the VA gender question. The choice not to include transgender as an option for VAs was due to the absense of biological sex in VA technologies. The other differences between the two questions was that participants could opt not to provide their gender, and could select that the gender of a VA "Does not matter."

Notably, for both questions, brief definitions were provided in parenthesis for a few of the genders, for example, "Agender (does not relate to the idea or experience of gender)" and "Bigender (both male and female at varying points)." These descriptors were included for participants who might not be familiar with these terms.

The responses to the question of what gender VAs should portray (Figure 6) starkly contrast the gender distribution of participants (Figure 7). For example, although no respondents described themselves as agender, 32.2% selected it as the gender that VA should portray. An additional two respondents mentioned agender with added explanations through the "Other" option in the survey. One respondent stated, "agender, but not in the way humans experience it, but more like the way that my lunch kit experiences gender, which is to say not at all." The popularity of the "Agender" response is likely because respondents understand that VAs, as AI technologies, are incapable of actually relating to the experience of gender and therefore should not portray a gender. The inclusion of the descriptor mentioned in the previous paragraph of agender as meaning "does not relate to the idea or experience of gender" may have also unintentionally suggested to participants that this was somehow the 'correct' answer.

Figure 6

Response to the Question, "What gender should a VA portray?"

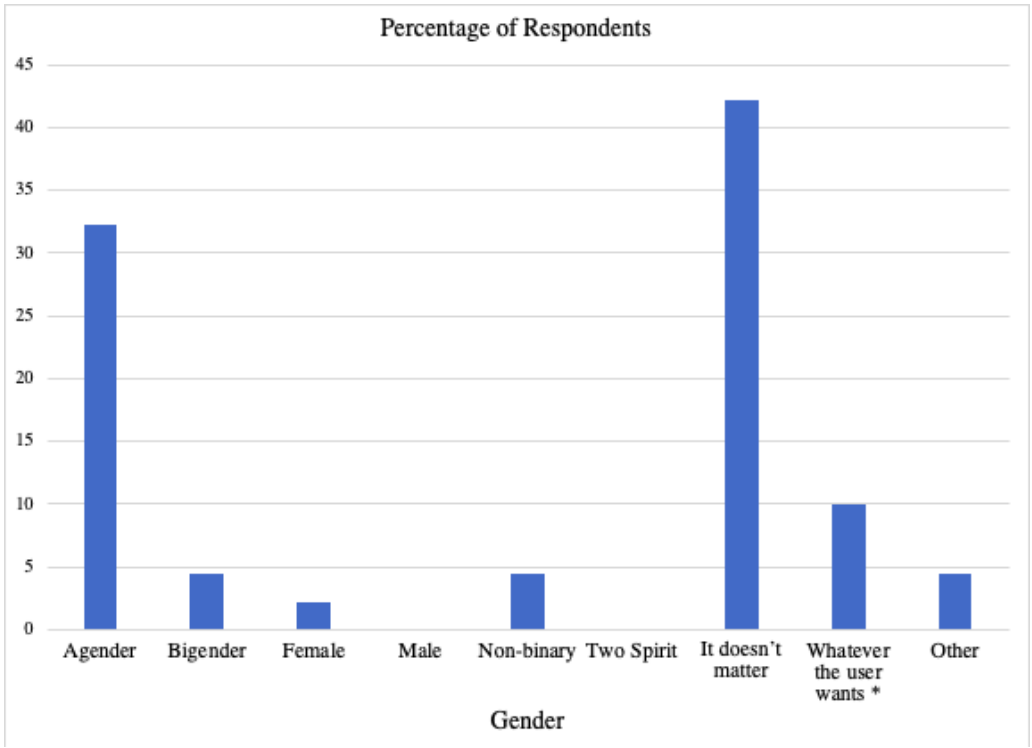

**Note:* The "Whatever the user wants" option was not actually provided for the question, but 10% of respondents mentioned that users should be able to customize the VA gender using the "Other" option. Because this number is significant it has been included as its own column. The remaining "Other" responses (4.4%) are represented in a separate column, 2.2% of these were some version of agender described above. A total of 14.4% of respondents used the other option to respond to this question in total.

Figure 7

Gender of Participants

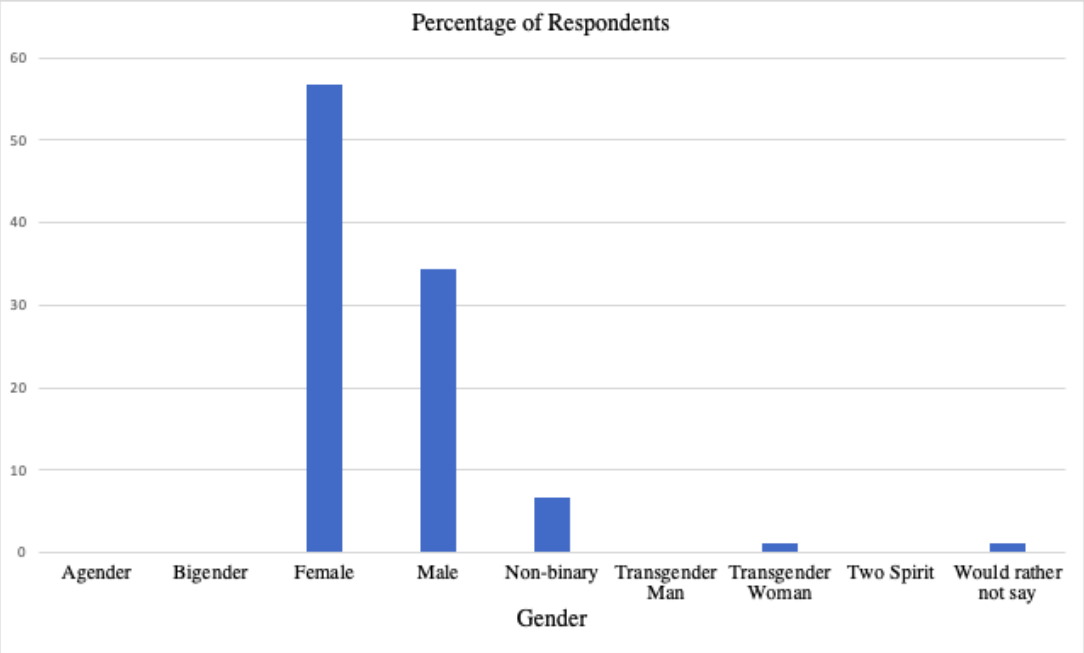

It is important to draw further attention to the phrasing of the question. Rather than asking participants what gender they would personally prefer a VA to portray, the question was phrased, "What gender should VAs portray?" The "should" here may have implied that respondents might consider ethical or moral implications when choosing an answer. It could have been illuminating to include an additional question asking what gender they would personally prefer their VA to portray for comparison. The results to the question (Figure 6) should therefore not necessarily be interpreted as representing participants personal preferences. One limitation of

these results is that it is not possible to confirm that every respondent interpreted this question the same way.

In total, 14.4% of respondents used the "Other" option to respond to this question using the text box provided. A significant number of respondents used this option to state that the gender should be customizable by the user. Clearly those respondents who felt that gender should be customizable were thinking of this question in terms of society generally rather than sharing their own personal preference for VA gender performance, or they would have simply selected the gender of their choice. An earlier draft of the survey included the option that gender could be selected by each user; however, this option was removed to encourage respondents to be more decisive in their answers.

Despite this attempt to encourage decisiveness, the most popular answer to the question of what gender a VA should portray was that it "Does not matter." 42.2% of respondents chose this answer and it is easy to imagine different reasons for its popularity. Potentially some respondents simply have no preference or do not feel that whether people interpret a voice as being typically female or typically male matters. Respondents could also feel uneasy about choosing a specific gender, wanting VAs to all portray different genders to better reflect human gender diversity.

Respondents are not alone in their resistance to committing to an assigned gender for a VAs, the companies themselves struggle with this. Neither Apple nor Google directly associates voices with genders anymore (they use numbers or colours to label the different options), and Amazon does not let other companies refer to Alexa as "she" or "her," even though they tend to themselves (Gartenburg, 2021). Interestingly, if you ask Alexa (at least in 2017) about gender, the response you would receive is that "gender identity is complex and personal, there is no definite way to say how many genders there are" (The Guardian, 2017).

Only 2.2% of respondents selected that VAs should portray either gender on the man-woman binary, despite the fact that 91.1% of respondents describe themselves as either male or female. This suggests that there might be interest among respondents for VAs to perform genders outside of the man-woman binary. The task of creating non-gender-conforming VA voices has been taken on by at least two different initiatives (Behr, 2020; Simon, 2019). The idea behind Q, a "genderless" voice created by a team of linguists, technologists, and sound designers, is to pressure the tech industry to acknowledging the complexity of gender (Simon, 2019).

According to linguist Kristina Hultgren, voices like the one used for Q challenge listeners and play with our natural tendency to put people into boxes (Simon, 2019). It is unclear if major tech companies will see value in incorporating voices that would push their consumers boundaries or if they would rather continue relying on more familiar archetypes and personas.

The Performance of Age

When asked what age VAs should portray, 58.9% of respondents stated that it did not matter. This is an even higher percentage than those who stated that gender portrayal did not matter (42.2%). As with gender, several respondents expressed that the age portrayed should be up to the user of the device.

If another question was included in the survey along the lines of, *Should AIs only portray children under 10?* it is unlikely that the percentage of respondents who said that age does not matter would remain so high. Presumably the people at Apple and Google would know better than to use child-like voices to give people directions or update them on how their stocks are doing. Perhaps, VAs portraying children would be so obviously unethical, that many respondents were not considering it as a factor when they selected that age does not matter. No respondents felt that VAs should portray an age under 20.

In total 30% of respondents expressed that neither the age nor the gender mattered. As discussed throughout this chapter, we tend to associate these technologies with human personal assistants or servants. It is interesting to consider whether people project popular conventions, attitudes, or standards for human employment onto VAs. In terms of the performance of age and gender, people may feel that they should be equal-opportunity VA 'hirers' and avoid sexist or ageist casting when it comes to VA characters. They might respond that age "Does not matter" because it shouldn't matter for people.

The second most popular answer for this question was that VAs should portray an age between 30-41 which was selected by 17.8%, followed by ages 41- 50, which was selected by 6.7%. These ages may have been selected because these are conventionally considered the key 'working years.' It might feel more appropriate for people to ask a VA portraying an age between 30-50 for assistance, rather than asking for help from a VA who sounds like it should be enjoying its retirement. This explanation does not account for the relatively low number of respondents (3.3%) who selected ages 20-31. Many personal assistants in popular culture are depicted as falling within this age range. To continue with speculation, slightly older voices may carry a

greater sense of authority or experience, which could be especially desirable in a technology that primarily answers questions. This could explain why respondents were drawn to the slightly older age range of 30-41 over voices that sounded like they were in their 20s.

It is difficult to always recognize the influence that social conditioning has on the associations we make based on age and gender. It is understandable that so many respondents expressed that the age and gender of VAs "Does not matter" because in a perfect world it probably shouldn't (with the exception, perhaps, of representing children). However, the reality is people do have culturally-informed associations that can be employed by VA developers to design technologies we will associate with the kind of person we would unquestioningly allow into our homes, cars, and workplaces.

Public Opinion on VA Personality

To gather opinions on which particular personality traits VAs should perform, participants were given a list of 28 human traits and asked to select all the traits that VAs should have. After responses were collected and the analysis process began, it became clear that some of the traits used for this question can be interpreted as describing the VA's functionality rather than portrayal of personality. Therefore the traits were divided into two categories: "task-related traits" and "personality-related traits." The traits "Clear," "Reliable" and "Helpful" were classified as task-related because these terms can be used to describe how successful a VA is at completing its primary technical functions, such as information retrieval and providing directions. The other traits that were included for this question can be interpreted as having more to do with the performance of personality and as such, have been classified as "personality-related traits."

Task-related Traits

The three task-related traits were the most popular answers (Figure 8). "Clear" was selected by 81.1% of respondents, "Helpful" by 68.9%, and "Reliable" by 44.4%. All three of these traits relate to how effective VAs are at completing tasks such as responding to queries, providing directions, taking notes, activating smart-home devices, etc.

Clear

It is understandable that this would be the human-like trait most frequently selected by respondents. Clear can refer both to how easy it is to understand the words a VA is saying and the directness of the responses. Obviously users would need to be able to hear and recognize the words a VA uses, but they also need answers that are direct and straightforward. A VA that

answered you in only riddles would be amusing at a party, but not helpful when providing directions.

Clarity is essential for VAs to effectively respond to queries and confirm that they have responded appropriately to commands. That said, how human-like does a voice need to be in order to be clear? Arguably, a certain degree of flow that resembles natural human conversation could assist with clarity. It is probably easier to understand a voice that sounds familiar rather than having to work to register meaning from clumsy, disjointed speech. Although, we have likely become more accustomed to automated voices over the last few decades and are therefore better practiced at understanding robotic-sounding speech.

Helpful

The popularity of "Helpful" might be related to the value of convenience discussed in Chapter 3. Interestingly 68.9% of respondents feel that VAs should be helpful, compared to 44.4% who actually described using VAs as being helpful earlier in the survey. It is unclear if this suggests that VAs are under-performing in the area of helpfulness.

Although "Helpful" has been listed as a task-related trait, somewhat synonymous with 'useful,' helpful can also be thought of as an attitude. According to Amazon, " Alexa lives in the cloud and is happy to help" (Amazon, n.d.a), and Apple describes Siri as "always learning how to be even more helpful" (Apple, n.d.).

Reliable

Inconsistent or fallible technology is never desirable. Which is likely why so many respondents stated that VAs should be reliable. For some, the word "reliable" may also carry connotations related to loyalty, like the reliable valet stereotype discussed in previous sections. The performance of loyalty will be discussed further in the final section of this chapter.

Figure 8

Percentage of Respondents who Selected Each Task-related Trait

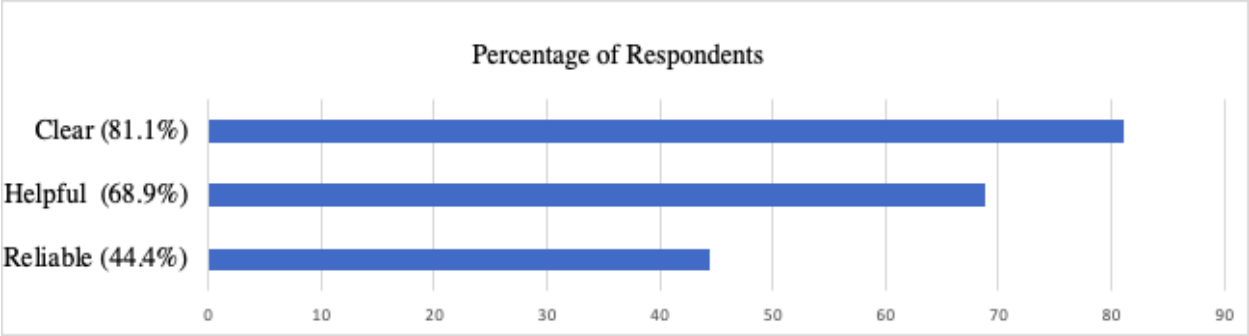

Personality-related Traits

Of the personality-related traits, "Polite" "Discrete," and "Friendly" were the most popular (Figure 9). This is potentially because these traits are commonly associated with good manners, at least by most western standards. They also could also be associated with generally accepted standards for good customer service. Returning to the idea of VAs as servants, all of these qualities would be desirable in the faithful valet character discussed previously or the stereotypical reliable secretary.

Figure 9

Percentage of Respondents who Selected Each Personality-related Trait (Note that the scale is different than Figure 8)

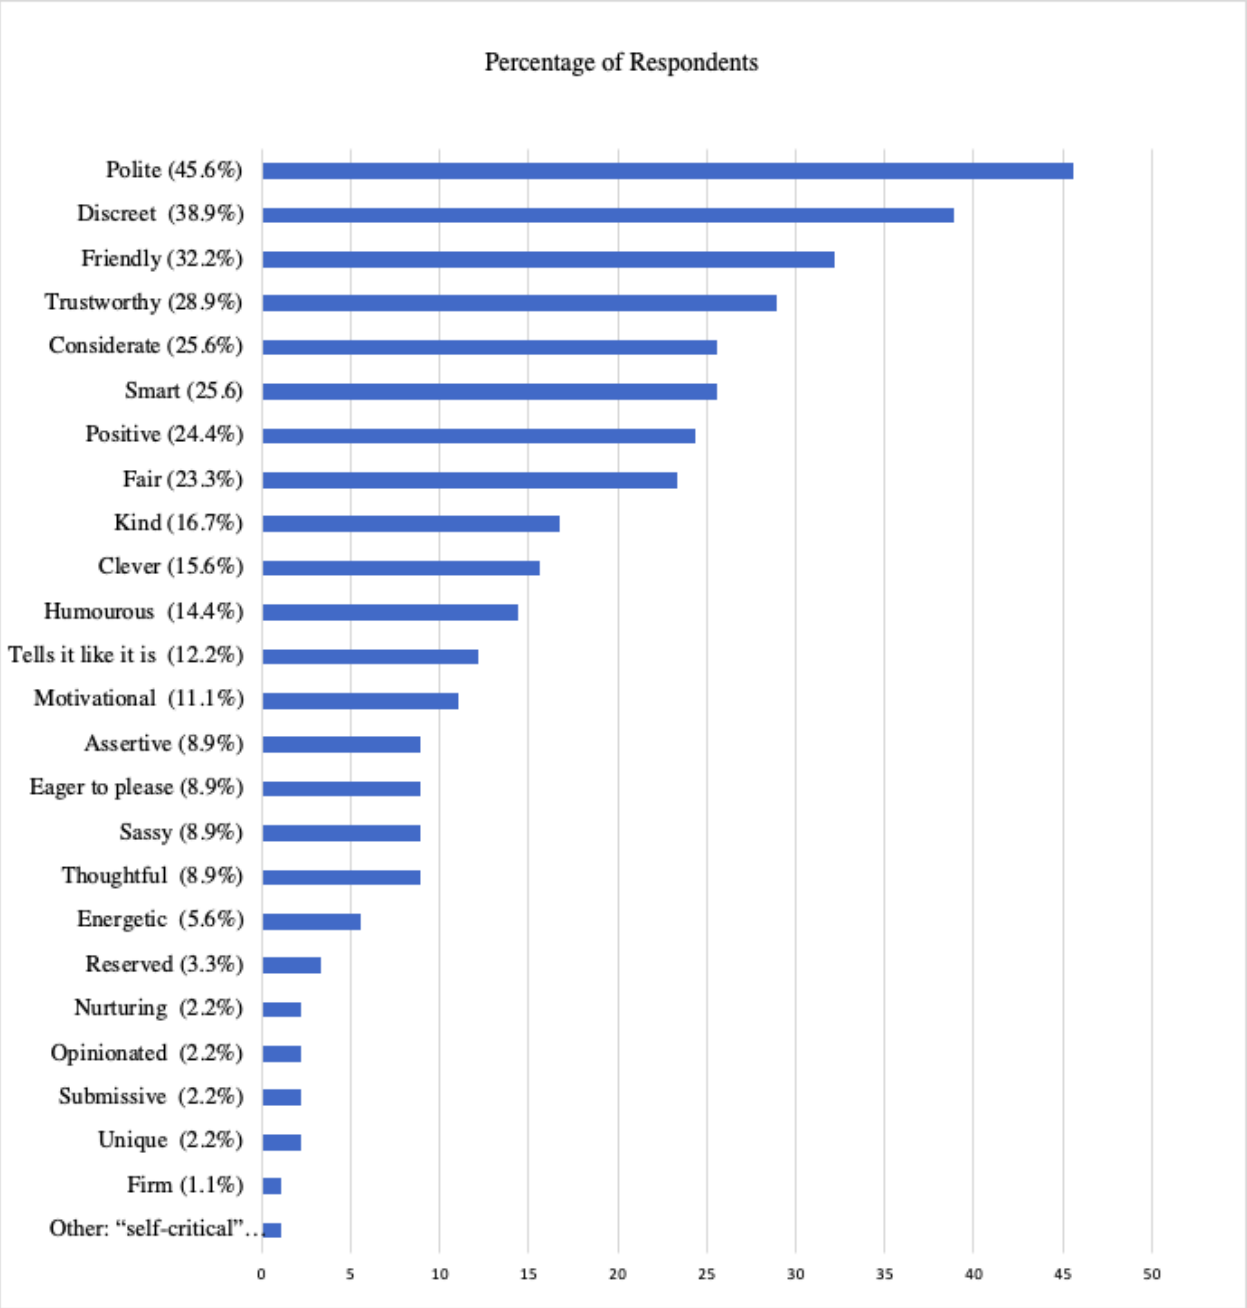

Politeness

The popularity of politeness is unsurprising. It is easy to assume that most people would not enjoy a rude, offensive voice yelling at them from their device. A study that examined reactions to AI in customer service roles, found that participants expected robots to be polite and conform to social norms towards them (Barnett et al., 2014). Interestingly, that research also found that people did not feel they had to reciprocate politeness toward the AI, and because of

this robots were expected to behave "in a super-human manner" maintaining politeness while being treated rudely (p. 1189). Presumably, this expectation that robots must conform to social norms, but needn't be treated like a human being itself, would apply to voice-based AI like VAs as well.

Discreet

Discrete was the second most popular personality-related trait. Discrete can mean showing discernment or good conduct and judgment, especially in speech (Merriam-Webster, n.d.b). By this definition, to be discrete is arguably a high-level human capability, involving subtlety and evaluation skills. Participants may also have been associating "Discrete" with being unobtrusive or even trustworthy. In any case, the popularity of this answer suggests that people want VA technologies to respond in a socially appropriate way, for example, a discrete VA would probably not interrupt the user.

Friendly

Friendly was the third most popular answer for personality-related traits. This is understandable given that friendly interactions are almost always preferable to unfriendly ones. Like with human-to-human interactions, when AI exhibits friendliness it may enhance a user's willingness to interact with it (Barnett et al., 2014). "Friendly" is related to "Polite" but perhaps with a bit of added warmth. Arguably exhibiting a "friendly" quality is shifting slightly toward a more human portrayal, by doing more than just the minimum to comply with social norms.

Trustworthy

"Trustworthy" was selected by 28.9% of respondents. Apple promotes Siri's ability to access personal information as a key feature on their website: "It's amazing how quickly Siri can find photos of your kids, locate your parked car, get you a ride, or unearth what you're looking for from massive numbers of files" (Apple, n.d.). With this access to personal data it is understandable that a significant portion of respondents want VAs to be trustworthy.

An argument could be made for categorizing "Trustworthy" as a task-related trait because the management of personal data is a technical aspect of VAs; however, what a VA does with data, does not necessarily impact how successfully they can respond to queries or commands which is why it has been classified as a personality-related trait for this research. When respondents express that they wish for VAs to be trustworthy, it may be that they really want VA *designers* to be trustworthy, rather than the technologies themselves.

Considerate

"Considerate" was the fifth most popular personality-related trait. This is an interesting personality trait because more than any of the other traits discussed so far, consideration has implications for machine learning. To some extent consideration implies that a VA can predict a user's needs. Apple's phrasing on their webpage devoted to Siri states, "Even when you don't ask, Siri works behind the scenes like a personal assistant" (Apple, n.d.). The anticipation of our needs and wants is a powerful quality that has great potential to endear VAs to us.

Consideration also shares some commonalities with another powerful human capability, empathy. When we show consideration for others and anticipate someone else's needs, we are to some extent, demonstrating understanding or recognition of each other. Although a VA cannot share our human experiences, it can demonstrate that it recognizes our needs based on patterns of behaviour.

Designing VAs so they come across as considerate or empathetic can be a powerful tool for VA developers. According to Pelau et al. (2021) human-like AI can be perceived as a threat to human identity, unless it can demonstrate empathy, "which changes the consumer's perspective" (p. 7). This is supported by other research that suggests that people tend to personify voice-activated smart technologies and look at them as sympathetic companions who have an interest in its user's physical and psychological well-being (Foehr & Germelmann, 2020).

Empathy and sympathy were not listed as traits that participants could select in the survey for this research, however traits like considerate, kind, trustworthy, and fair are similar characteristics that were identified by many of the respondents as being desirable for VAs. Further research could explore the relationship between VA ability to predict behaviour, perceived empathy, and adoption of VA technologies.

Humorous

Unlike the personality-related traits listed above, "Humourous" was not an especially popular trait. Interestingly, only 14.4% of respondents think that VAs should be humorous, but 37.8% stated that they used VAs for "Fun (asking funny questions etc.)". There are different possible explanations for this. It could be that users's fun derives from playing games by accessing applications through VAs rather than from the interaction itself. Another explanation could be that users prefer to get amusement by messing with VAs, and are not especially interested in a VA stand-up comedian. There may be commonalities with Barnett et al.'s 2014

research around politeness that was discussed above. Potentially, users expect to be able to make fun of VAs, while not having the VAs make jokes themselves. This may also tie into the servant character discussed previously (it wouldn't be a VA's 'place' to crack jokes). Further research would be needed to explore these theories.

Other Personality Traits

Other popular choices for personality-based traits were smart, positive, fair, kind, and clever. Notably, even among those who had previously stated that a VA should only be human-like enough that you can communicate with it verbally but not more, there were 21 respondents who selected personality traits, such as humorous, sassy, considerate, assertive, motivational, eager to please, or thoughtful, as traits VAs should have. This might suggest that even though these respondents do not want to feel like they are having a conversation with a real person (another option for that question) they wouldn't mind a conversation that is a bit more interesting. Maybe an obviously robotic voice saying something motivational or sassy would be amusing, for example.

"No Human-like Traits"

6.7% of participants said that VAs should have no human-like traits. It is hard to imagine a system that communicates through natural language technology having absolutely no human-like qualities. As mentioned in the introduction to this chapter, the human voice carries a lot of information for us to interpret. People also tend to anthropomorphize AI (Scheuz, 2021; Foehr & Germelmann, 2020). The respondents who believe that VAs should have no human-like traits may recognize that the portrayal of human-like qualities, although unavoidable, is inherently deceptive. An AI does not operate like a human brain, yet it communicates the way humans do.

Should AI Identify Itself?

In 2018 Google gave a demonstration of a new, natural language AI technology called Google Duplex. In the demonstration the technology booked a hair appointment over the phone at a (presumably real) hair salon. Google Duplex did an "uncannily good job of asking the right questions, pausing in the right places, and even throwing in the odd 'mmhmm' for realism" (Vincent, 2018). This demonstration was met with controversy as people voiced their concern that AI could be used to intentionally trick people into thinking they were talking with a human

(Statt, 2018). Following the outcry, Google rapidly made assurances that Duplex would identify itself as AI; however, this self-identification was not mentioned in the original demo.

Most of the respondents (78.9%) believe that a VA should have to identify itself as AI. 12.2% of respondents said that whether an AI needs to identify itself depends on the situation. Respondents who chose "It depends" were encouraged to explain their reasoning in a text-box. Of those who selected "It depends": four respondents explained that a VA should identify itself when it is not obvious that you are talking to an AI; three respondents felt that whether an AI should identify itself had to do with the complexity of the information or how sensitive it is (e.g. if an AI is providing health information it should identify itself, but for weather information it wouldn't need to); and three other respondents felt it depended on the direction of the interaction (e.g. if an AI initiated the interaction, or was soliciting something from them, they would want to know they were dealing with an AI).

Although the overwhelming majority of respondents want to know when they are talking to AI, whether that is because it is obvious or explicitly communicated, there were those who felt a VA did not have to identify itself as AI. At the moment, VA technologies are easily recognizable as AI because their natural language capabilities are not perfect. For the 8.9% who stated that VAs did not have to identify themselves as AI, it is not possible to tell whether they feel it is not necessary because it is obvious or if they simply do not mind.

Deception, Loyalty and Manipulation

When asked if they had any other thoughts on Virtual Assistants they would like to mention, one respondent wrote:

> "I think I'd hate them more if they weren't so cute. Sure I get frustrated when they don't work properly, but I can't get too angry at something that talks like a human and has been put on this planet to help me. I know it's just a computer but I still humanize them."

It is very striking that this respondent describes a VA as having been put on the planet to help them. This arguably speaks to how successful companies have been at evoking the assistant/servant/secretary stereotype into the voices and 'personalities' of VAs.

When we think of a VA as trusty valet, or on-the-ball secretary there is also an implied loyalty. According to Amazon, " Alexa wants to make your life easier, more meaningful, and more fun by helping you voice control your world" (Amazon, n.d.a). This phrasing suggests that

Alexa is capable of "wants," cares for users, and has their best interest at heart. It seems that this statement is written to evoke evokes feelings of empowerment, loyalty, and trust.

There are ethical questions that emerge when it comes to the portrayal of human capabilities like kindness, empathy, and trustworthiness by VAs. For example, are VAs being designed to be empathetic, or to *sound* empathetic? The performance or suggestion of empathy, without the capacity for it, could be considered deception. However, deception is often forgivable if it is for the right reasons. Some argue that the shared ability to deceive might even endear humans to AI, as it did with the robot R2D2 in the original *Star Wars* films (Isaac & Bridewell, 2017). Future research could explore public opinion around deception by AI further. In particular, what people consider acceptable deception by AI and how that compares with acceptable deceptions by humans.

Arguably, VAs impersonating human behaviours, emotions, and personality is not considered deception but rather manipulation. When we initiate an interaction with a VA we understand that we are communicating with an AI and are therefore not being completely deceived about its capabilities and limitations. That said, even though we know we are not talking to a real person we might still be influenced by the human-like qualities performed by a VA. For example, we can still form emotional attachments. Dr. Cynthia Breazeal, the creator of Kismet, a robot designed to recognize and imitate human emotions, developed a "maternal connection" to the AI and experienced a sharp sense of loss when she had to leave the robot behind (Scheutz, 2012, p. 213). This example demonstrates that even when you thoroughly understand the workings of an AI you can still become very emotionally invested.

Summary

The goal of this chapter was to explore the public opinion data collected in the survey about if and how VAs should perform human-like characteristics. Looking at the responses it is clear that in many ways public opinion about which personality traits VAs should have aligns with the way that VA characters are designed to come across. The traits that respondents most frequently felt VAs should have were clear, helpful, polite, reliable, discrete and friendly. All these traits can be recognized in the way that VA's currently attempt to portray personality and the way companies promote these technologies. Many of these traits also align with commonly accepted social standards for good manners and customer service.

In terms of the portrayal of gender, very few respondents, 2.2%, selected a gender on the man-women binary (in both cases female) as being the gender that VAs portray. This stands in sharp contrast to current practices of major VA developers who consistently use voices that can be identified as either male-sounding or female-sounding. Notably, in almost every case a female-sounding voice has been the default.. Both in terms of the portrayal of gender and the portrayal of age, high percentages of respondents felt that the age or gender a VA performed did not matter. Potential explanations for this are described in previous sections.

An area where respondent views sharply contrasted the trajectory of VA development was the question of how human-like VAs should be. The majority of respondents feel that VAs should be minimally human-like which stands in contrast to ongoing efforts by companies like Google to make natural language technology that sounds increasingly human-like.

Chapter 7: Conclusion

This research aimed to build off of public opinion data to explore ethical concerns and the anthropomorphization of AI in Virtual Assistant (VA) technologies. Data was collected through a voluntary online survey using non-probability sampling. Given that the survey was distributed using convenience sampling the results of the survey should not be interpreted as generalizable findings but are rather meant to provide indications of potential trends in public opinion around VAs and reveal opportunities for future research.

Summary of Findings

The results of the survey suggest that there are multiple issues with VA technologies that produce significant levels of concern for respondents. Concerns related to corporate manipulation, surveillance and lack of transparency, as well as concerns related to political manipulation or increased political polarization, were the issues where participants most frequently expressed high levels of concern. Issues related to autonomy and dependence less frequently elicited intense concern. This might indicate that people are less concerned about issues related to personal decision-making and behaviour and have greater concerns about large-scale issues like mass surveillance by corporations, than they are about risks to themselves as individuals.

Responses also suggested that the level of concern around VA technologies is relatively similar to the level of concern participants had for technology in general. Additionally, respondents overwhelmingly had low levels of trust that technology companies that make VAs are making sure their technologies do not negatively impact people, society, or the environment. Findings also suggest that respondents believe multiple parties should be responsible for ensuring that VA technology does not cause harm. There may, therefore, be public support for a collaborative approach to the regulations of VA development which might involve different governments, organizations, individuals, and the companies themselves.

In many ways respondent opinions about which personality traits VAs should have, align with the way that VA's "personalities" are currently designed to come across. The human-like traits that respondents most frequently believed that VAs should exhibit were clear, helpful, polite, reliable, discrete and friendly. Notably many of these traits align with commonly accepted social standards for good manners and customer service.

Although many of the personality traits deemed desirable by respondents can arguably be perceived in VAs on the market today, there were other instances where the opinions represented in the data contrasts the ways that VAs are currently designed. For example, very few respondents (2.2%) expressed that VAs should portray a gender on the man-women binary (in both cases female). This contrasts the fact that VA developers consistently use voices that can be identified as either male-sounding or female-sounding. That said, it should be noted that many respondents expressed that the gender performed by a VA did not matter. Another instance of respondent opinion sharply contrasting VA development was that the majority of respondents stated that VAs should be human-like enough that you can communicate with them verbally but not more. This contrasts ongoing efforts by companies like Google to make natural language technology that sounds increasingly human-like.

Areas for Future Research

The analysis of survey responses has led to the identification of several areas where future research is needed. The following section summarizes the recommended areas for future research discussed throughout the last three chapters.

Public Opinion and VA Design Trajectories

As discussed in the previous chapter, the majority of respondents believed that VAs should be minimally human-like, which contrasts the path that VA technology is currently on. VAs and other Social AI technologies may be completely unavoidable in the future, and yet very few people have a say in how these technologies are designed. Further research is needed to compare public opinion to technical development trajectories in greater detail. This could involve examining if and how consultation or public opinion research could be used more effectively to critically examine the social impacts and ethical implications of these technologies. There may also be potential for this kind of research to be used to increase public awareness about new technologies or to inform industry standards and practices.

Comparing Different VAs

One of the limitations of this research is that it lumps all VAs together and gathers opinions on these technologies generally. When asked if they had any other thoughts on VAs they would like to mention, one respondent wrote:

> "I view VA's as being different depending on the company that made it. I.e. Amazon Alexa is made to sell more things. Google Assistant is made to sell more ads. Apple Siri

> is made to help the user. While this may not be 100% true, Apple's marketing around privacy certainly makes me feel better about using its VA vs the other companies. Personally, I'd never use Alexa.."

The associations that users have with the companies creating VAs will impact the level of trust they have for these technologies. A priority for future research will be to look closer at the differences in public perception between VAs developed by different companies.

Perceived Empathy and VA Adoption

The performance of empathy by anthropomorphic AI has the potential to endear VAs to their users despite the fact that as AI, VAs lack the capacity for empathy. Future research should explore the relationship between VA design, perceived empathy, trust and adoption of VA technologies in greater detail. There may also be implications here for other Social AI technologies.

Acceptable Deception

Future research should examine public opinion around deception by AI further. This should include examining when deception is acceptable in humans and whether that translates to VA technologies. Additionally, exploration is needed to determine if companies suggesting that VAs have human-like capabilities (either directly or indirectly) through their promotion of VAs, is problematic or acceptable.

Abuse of VAs

As discussed in the previous chapter, VAs might be expected to be polite even when treated rudely, as was found with research involving other kinds of AIs (Barnett et al., 2014). Potentially there is also an expectation that users want to be able to make fun of VAs, while not having VAs make jokes themselves. This might suggest that there is satisfaction that comes from treating a human-like object, inhumanly. Further research should explore this in greater detail including examining other implications of this either on individuals or society.

Optimism Bias and VAs

Cho et al. (2010) observed that people can display a strong optimism bias about online privacy risks (p. 994). Future research should explore whether users of VAs perceive themselves as less vulnerable than others to risks such as threats to autonomy, technology dependence, or the threat of VAs taking human jobs.

Awareness and Concern Around Data Collection

The survey findings seem to suggest that surveillance is a significant concern regardless of how aware respondents feel they are about the collection and use of data. Future research should use a larger sample size composed of both experts and those with no familiarity of how data is collected to explore the relationship between awareness and concern in greater detail. One example would be comparing people's own perceived dependence on technology, with how they would describe the dependency on technology of others, including both close contacts and society in general.

Fun and VAs

Future research should explore the relationship between fun and adoption of VAs, including whether enjoyment of these technologies comes from the novelty of interacting with an AI or if it is part of a lasting positive user experience that motivates regular use of VAs. Further research could also explore the relationship between enjoyment of VAs and trust in technology companies.

Social Hesitancy

People may feel less hesitant to interact with a VA than another person for many reasons. Further research should explore the appeal of interacting with a VA over interacting with another human being and the potential implications of replacing human interactions with AI interactions.

Perceived Expertise and VAs

A significant portion of respondents would rather get directions from a VA than another person, this may be because we perceive VAs to be experts in directions given their access to map and GPS data. Future research could explore if VAs programmed with specialized information, such as medical information, would be trusted to respond appropriately to related queries to a greater extent than the popular VA options currently on the market. This research might reveal if VAs might be adopted for a wide range of different information needs if they are trained with reliable data rather than just drawing from the internet for responses.

Risk and Complexity

How complex a question or command is, and the degree of risk associated with an incorrect response, could be key factors in whether people trust a VA with a particular task. As these technologies become more popular and the range of tasks VAs are capable of expands, future research should explore the factors that contribute to how we determine which tasks should be delegated to VA technologies.

Final Thoughts

With the advancement of natural language technology, it will not be long before VAs can accurately mimic human interactions. In spite of the fact that VAs will not have the capacity for human-like experiences such as empathy, they will become more capable of imitating those experiences. Further research on the ethical implications of human-like AI is urgently needed. Currently, the decisions about how these technologies should be designed are in the hands of a very small number of people who are, in large part, embedded in the profit-motivated surveillance capitalist system.

Because these technologies have the potential to become integrated into many aspects of society, ongoing research should collect input not only from AI experts, companies, and academics but from members of the public, including both those who are users of technologies like VAs and those who are not. Going forward, gathering and incorporating input from a diverse cross-section of society is essential to align the design and regulation of Social AI, with the priorities of those who will increasingly encounter these technologies in the future. If this research does nothing else it demonstrates that a great deal of information can be discovered by engaging people in thinking critically about AI design.

References

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science, 347*(6221), 509-514. DOI: 10.1126/science.aaa1465

Algorithm Watch. (n.d.). AI ethics guidelines global inventory. Retrieved November 9, 2021, from https:// inventory.algorithmwatch.org

Amazon (n.d.a). Alexa featured. Retrieved November 15, 2021, from https://www.amazon.com/b?node=21576558011

Amazon (n.d.b). Designed to protect your privacy. Retrieved December 18, 2021, from https://www.amazon.com/b/?node=19149155011

Amazon (n.d.c). What is Alexa? Retrieved November 29, 2021, from https://www.developer.amazon.com/en-US/alexa.

Apple (2021a, April 26). iOS 14.5 delivers unlock iPhone with Apple watch, more diverse Siri voice options, and new privacy controls. https://www.apple.com/newsroom/2021/04/ios-14-5-offers-unlock-iphone-with-apple-watch-diverse-siri-voices-and-more/

Apple (2021b, October 27). Apple Privacy Policy. https://www.apple.com/legal/privacy/en-ww/

Apple (n.d.). *Siri does more than ever. Even before you ask.* Retrieved November 14, 2021, from https://www.apple.com/siri/

Anic, I.-D., Budak, J., Rajh, E., Recher, V., Skare,V., & Skrinjaric, B. (2019). Extended model of online privacy concern: What drives consumers' decisions? *Online Information Review, 43*(5), 799-817. DOI: 10.1108/OIR-10-2017-0281

Barnett, W., Foos, A., Gruber, T., Keeling, D., Keeling, K., & Nasr, L. (2014). Consumer perceptions of interactive service robots: A value dominant logic perspective. *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, 1134-1139. DIO:10.1109/ROMAN.2014.6926404

Becker, M. (2019). Privacy in the digital age: Comparing and contrasting individual versus social approaches towards privacy. *Ethics and Information Technology, 21*(4), 307-317. DOI: 10.1007/s10676-019-09508-z

Behr, M. (2020, December 17). Meet Sam: The non-binary voice for a digital assistant. *Digit.* https://www.digit.fyi/meet-sam-the-non-binary-voice-for-a-digital-assistant/

Bond, R. M., Fariss, C. J., Jones, J. J., Kramers, A.D.I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature, 489,* 295–298. DOI: 10.1038/nature11421

Borau, S., Otterbring, T., Laporte, S., & Wamba, S. F. (2021). The most human bot: Female gendering increases humanness perception of bots and acceptance of AI. *Psychology & Marketing, 38*, 1052–1068. DOI: 10.1002/mar.21480

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science, 358*(6370), 1530-1534. DOI:10.1126/science.aap806

Büchi, M., Fosch-Villaronga, E., Lutz, C., Tamò-Larrieux, A., Velidi, S., & Viljoen, S. (2020). The chilling effects of algorithmic profiling: Mapping the issues. *Computer Law & Security Review: The International Journal of Technology Law and Practice*, *36*. DOI: 10.1016/j.clsr.2019.105367

Cho, H., Lee, J. S., & Chung, S. (2010). Optimistic bias about online privacy risks: Testing the moderating effects of perceived controllability and prior experience. *Computers in Human Behaviour, 26*(5), 987-995. DOI: 10.1016/j.chb.2010.02.012

Christl, W. (2017). Corporate surveillance in everyday life: How companies collect, combine, analyze, trade and use personal data on billions. *Cracked lab – Institute for critical digital culture.* https://crackedlabs.org/dl/CrackedLabs_Christl_CorporateSurveillance.pdf.

Christensen, H., & Eriksen, S. (2021). Is your voice enough, Alexa? Assessing the role of digital assistant personality, modality, and product involvement on consumer evaluations. [Masters Thesis, BI Norwegian Business School] https://hdl.handle.net/11250/2827760

Coleman, J.S.M., (2018). Descriptive statistics. In B. B.Frey (Ed.), *The SAGE encyclopedia of educational research, measurement, and evaluation* (pp .488-492). DOI: 10.4135/9781506326139

Collins, E., & Ghahramani, Z. (2021, May 18). LaMDA: Our breakthrough conversation technology. *Google: The Keyword*. https://blog.google/technology/ai/lamda/

Crawford, K. (2021). *The Atlas of AI : Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press. www-jstor-org.login.ezproxy.library.ualberta.ca/stable/j.ctv1ghv45t

Debatin, B., Lovejoy, J. P., Horn A-K., & Hughes, B.N. (2009). Facebook and online privacy: Attitudes, behaviors, and unintended consequences. *Journal of Computer-Mediated Communication, 15*(1), 83–108. https://doi.org/10.1111/j.1083-6101.2009.01494.x

Deloitte (2018). Ethics in the age of technological disruption: A discussion paper for the 2018 true north conference. https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/deloitte-analytics/ca-en-analytics-FCC-true-north-aoda.PDF

Dubois, E., & Blank, G. (2018). The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information Communication & Society*, *21*(5), 729–745. https://doi-org.login.ezproxy.library.ualberta.ca/10.1080/1369118X.2018.1428656

Etherington, D. (2014, November 6). Amazon Echo is a $199 connected speaker packing an always-on Siri-style assistant. *TechCrunch.* https://techcrunch.com/2014/11/06/amazon-echo/

E-Marketer (2019). US voice assistant users 2019: Insider intelligence trends, forecasts & statistics. https://www-emarketer-com.login.ezproxy.library.ualberta.ca/content/us-voice-assistant-users-20

EU High Level Expert Group on AI (2019). Ethics guidelines for trustworthy AI. *Brussels: European Commission.* https://ec-europa-eu.login.ezproxy.library.ualberta.ca/futurium/en/ai-alliance-consultation/guidelines#Top

Fenwick, M. D., Kaal W. A., & Vermeulen, E. P. M. (2017). Regulation tomorrow: What happens when technology is faster than the law? *American University Business Law Review, 6*(3). http://digitalcommons.wcl.american.edu/aublr/vol6/iss3/1

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *The Berkman Klein Center for Internet & Society at Harvard University Publication Series*. http://dx.doi.org/10.2139/ssrn.3518482

Floridi, L. (2016). Faultless responsibility: On the nature and allocation of moral responsibility for distributed moral actions. *Philosophical Transactions of the Royal Society A:*

Mathematical, Physical and Engineering Sciences, 374, 1-13. DOI: 10.1098/rsta.2016.0112

Foehr, J., & Germelmann, C.C. (2020). Alexa, can I trust you? Exploring consumer paths to trust in smart voice-interaction technologies. *Journal of the Association for Consumer Research, 5*(2), 181-205. DOI: 10.1086/707731

Friedman, B., Kahn, P. H. Jr., Borning, A., & Huldtgren, A. (2013). Value sensitive design and information systems. In N. Doorn, D. Schuurbiers & I. van de Poel, M. E. Gorman (Eds.), *Early engagement and new technologies: Opening up the laboratory* (pp. 55–95). Dordrecht: Springer.

Frischmann, B. M., & Selinger, E. (2018). *Re-engineering humanity*. Cambridge University Press. DOI: :10.1017/9781316544846.004

Gartenburg, C. (2021, July 21). Alexa finally gets a masculine-sounding voice option: A long-overdue addition for Amazon's assistant. *The Verge.* https://www.theverge.com/2021/7/21/22587130/alexa-masculine-voice-gender-ziggy-amazon-smart-assistant

Golossenko, A., Gopalakrishna Pillai, K., & Aroen, L. (2020). Seeing brands as humans: Development and validation of a brand anthropomorphism scale. *International Journal of Research in Marketing 37*(4). 737-755. https://doi.org/10.1016/j.ijresmar.2020.02.007

Google (n.d.). Artificial intelligence at Google: Our principles. Retrieved December 18 2021, from: https://ai.google/principles/

The Guardian (2017, December 12). Hey Alexa is it true your a lefty feminist? *The Guardian* [online] https://www.theguardian.com/technology/shortcuts/2017/dec/12/hey-alexa-is-it-true-your-e-a-lefty-feminist

Gruber, T. (2017). How AI can enhance our memory, work and social lives. *TEDTalks.* https://fod-infobase-com.login.ezproxy.library.ualberta.ca/p_ViewVideo.aspx?xtid=160835

Heinrichs, J.-H. (2021). Why digital assistants need your information to support your autonomy. *Philosophy & Technology, 1*. DOI: 10.1007/s13347-021-00481-4

Hern, A. (2019, July 26). Apple contractors 'regularly hear confidential details' on Siri recordings. *The Guardian.*

https://www.theguardian.com/technology/2019/jul/26/apple-contractors-regularly-hear-confidential-details-on-siri-recordings

Hoa, K. (2021, May 20). The race to understand the exhilarating, dangerous world of language AI. *MIT Technology Review* [online]. https://www.technologyreview.com/2021/05/20/1025135/ai-large-language-models-bigscience-project/

IBM (n.d.). IBM Shoebox. Retrieved December 18, 2021, from https://www.ibm.com/ibm/history/exhibits/specialprod1/specialprod1_7.html

IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems (2019). Section 2: When systems care. In *Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems*. https://engagestandards.ieee.org/rs/211-FYL-955/images/EAD1e.pdf?mkt_tok=eyJpIjoiWmprek1tUmxZbU14T1RnMCIsInQiOiJTMVFuV2JMUlBLYUp4ekh4MXZcL2J2c1dTMnZPRlZqSkxUdWJEMHA1b1E5SGVSTlZxNWFLclFJdjRUMHlcL3VWcmp1amJManFQNmxmUFBrZWRnZ25PMzFjQWM3U2hYcmNmemt5U1VVWDNxTFwva2VEN0d1a3I0VnRUXC80Rm94bm9GNWcifQ%3D%3D

Isaac, A.M.C., & Bridewell, W. (2017). White lies on silver tongues: Why robots need to deceive (and how). In P. Lin, R. Jenkins & K. Abney (Eds.), *Robot ethics 2.0: From autonomous cars to artificial intelligence* (2nd ed., pp. 157- 172). New York, NY: Oxford University Press.

Jacobs, N., & Huldtgren, A. (2018). Why value sensitive design needs ethical commitments. *Ethics and Information Technology, 23,* 23-36. https://doi.org/10.1007/s10676-018-9467-3

Juang, B.H., & Rabiner, L. R. (2005). Automatic speech recognition – A brief history of the technology development*. Research Gate.* www.researchgate.net/publication/249888949_Automatic_Speech_Recognition_-_A_Brief_History_of_the_Technology_Developmenthttps://www.researchgate.net/publication/249888949_Automatic_Speech_Recognition_-_A_Brief_History_of_the_Technology_Development

Ki, C.-W. (Chloe), Cho, E., & Lee, J.-E. (2020). Can an intelligent personal assistant (IPA) be your friend? Para-friendship development mechanism between IPAs and their users. *Computers in Human Behavior*, *111*. DOI: 10.1016/j.chb.2020.106412

Kim, S., & McGill, A. L. (2011). Gaming with Mr. Slot or gaming the slot machine? Power, anthropomorphism, and risk perception. *Journal of Consumer Research, 38*(1), 94–107. DOI: 10.1086/658148

King, O. C. (2020). Presumptuous aim attribution, conformity, and the ethics of artificial social cognition. *Ethics and Information Technology, 22*(1), 25-37. DOI:10.1007/s10676-019-09512-3

Kitchens, B., Johnson, S. L., & Gray, P. (2020). Understanding echo chambers and filter bubbles: The impact of social media on diversification and partisan shifts in news consumption. *MIS Quaterly, 44*(4), 1619–1649. DOI: 10.25300/MISQ/2020/16371

Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers and Security, 64*, 122-134. DOI: 10.1016/j.cose.2015.07.002

Langer, M., & König, C. J. (2018). Introducing and testing the creepiness of situation scale (CRoSS). *Frontiers in Psychology*, *9*. DOI: 10.3389/fpsyg.2018.02220

Lara, F. (2021). Why a virtual assistant for moral enhancement when we could have a Socrates? *Science and Engineering Ethics*, *27*(4). DOI: 10.1007/s11948-021-00318-5

Lee, H., Park, H., & Kim, J. (2013). Why do people share their context information on Social Network Services? A qualitative study and an experimental study on users' behavior of balancing perceived benefit and risk. *International Journal of Human-Computer Studies, 71*(9), 862–77. DOI: 10.1016/j.ijhcs.2013.01.005

Lopatovska, I., Rink, K., Knight, I., Raines, K., Cosenza, K., Williams, H., Sorsche, P., Hirsch, D., Li, Q., & Martinez, A. (2019). Talk to me: Exploring user interactions with the Amazon Alexa. *Journal of Librarianship & Information Science*, *51*(4), 984–997. DOI: 10.1177/0961000618759414

Easwara Moorthy, A., & Vu, K.-P. L. (2015). Privacy concerns for use of voice activated personal assistant in the public space. *International Journal of Human-Computer Interaction, 31*(4), 307-335. DOI: 10.1080/10447318.2014.986642

Mawby, J., Foster, A., & Ellis, D. (2015). Everyday life information seeking behaviour in relation to the environment: Disposable information? *Library Review (64)*6/7, 468-479. DOI 10.1108/LR-10-2014-0120

McTear, M. (2020). *Conversational AI : dialogue systems, conversational agents, and chatbots*. Morgan & Claypool Publishers. https://search-ebscohost-com.login.ezproxy.library.ualberta.ca/login.aspx?direct=true&db=cat03710a&AN=alb.9365465&site=eds-live&scope=site

Merriam-Webster. (n.d.a). Manipulate. In *Merriam-Webster.com dictionary.* Retrieved November 8, 2021, from https://www.merriam-webster.com/dictionary/manipulate

Merriam-Webster. (n.d.b). Discrete. In *Merriam-Webster.com dictionary.* Retrieved November 8, 2021, from https://www.merriam-webster.com/dictionary/discreet

Microsoft (n.d.). Responsible AI. Retrieved December 18, 2021, from https://www.microsoft.com/en-us/ai/responsible-ai?activetab=pivot1%3aprimaryr6

Mori, M. (1970). The uncanny valley. In J. A. Weinstock (Ed.), *The Monster Theory Reader* [2020] (pp. 89–94). University of Minnesota Press.

Morley, J., Elhalal, A., Garcia, F., Kinsey, L., Mokander, J., & Floridi, L. (2021). Ethics as a service: A pragmatic operationalisation of AI Ethics. *Minds and Machines Minds, 31,* 239–256. DOI: 10.1007/s11023-021-09563-w

Moss, E., & Melcalf, J. (2019, November 14) The ethical dilemma at the heart of big tech companies. *Harvard Business Review.* https://hbr.org/2019/11/the-ethical-dilemma-at-the-heart-of-big-tech-companies

Muro, M., Whitin, J., & Maxim, R. (2019). What jobs are affected by AI? Better-paid, Better-educated workers face the most exposure. *Brookings.* https://www.brookings.edu/wp-content/uploads/2019/11/2019.11.20_BrookingsMetro_What-jobs-are-affected-by-AI_Report_Muro-Whiton-Maxim.pdf

Nguyen, A & Vu, H. (2019).Testing popular news discourse on the "echo chamber" effect: Does political polarisation occur among those relying on social media as their primary politics news source? *First Monday, 24*(6). DOI: 10.5210/fm.v24i6.9632

Ojeda, C. (2019). The political responses of virtual assistants. *Social Science Computer Review 39*(5), 884-902. DOI: 10.1177/0894439319886844

O'Reilly Media (2020, October 1). How data privacy leader Apple found itself in a data ethics catastrophe. *Medium.* https://medium.com/oreillymedia/how-data-privacy-leader-apple-found-itself-in-a-data-ethics-catastrophe-47afdcec77af

Pelau, C., Dabijia, D.-C., & Ene, I. (2021).What makes an AI device human-like? The role of interaction quality, empathy and perceived psychological anthropomorphic characteristics in the acceptance of artificial intelligence in the service industry. *Computers in Human Behaviour, 122.* DOI: 10.1016/j.chb.2021.106855

Pinola, M. (2011, November 2). Speech recognition through the decades: How we ended up with Siri. *PC World.* https://www.pcworld.com/article/477914/speech_recognition_through_the_decades_how_we_ended_up_with_siri.html

Price, M. S. (2020). Internet privacy, technology, and personal information. *Ethics and Information Technology, 22*(2), 163-173. DOI: 10.1007%2Fs10676-019-09525-y

Pirhonen, J., Melkas, H., Laitinen, A. & Pekkarinen, S. (2020). Could robots strengthen the sense of autonomy of older people residing in assisted living facilities?—A future-oriented study. *Ethics and Information Technology, 22*(2), 151-162. DOI:10.1007/s10676-019-09524-z

Russell, S., & Norvig, P. (2013). *Artificial intelligence: A modern approach* (3rd edition). Pearson.

Scheutz, M. (2012). The inherent dangers of unidirectional emotional bonds between humans and social robots. In P. Lin, K. Abney, & G.A. Bekey (Eds.), *Robot ethics: The ethical and social implications of robotics* (1st ed., pp.205-222). Cambridge, Mass.: MIT Press. https://ebookcentral.proquest.com/lib/ualberta/reader.action?docID=3339351&ppg=216

Sharif, K., & Tenbergen, B. (2020). Smart home voice assistants: A literature survey of user orivacy and security vulnerabilities. *Complex Systems Informatics and Modeling Quarterly*, *24*, 15–30. DOI: 10.7250/csimq.2020-24.02

Sherman, C. (2003, October 8). What's in a (search-engine's) name? *Search Engine Watch.* https://web.archive.org/web/20150103053032/http://searchenginewatch.com/sew/news/2063925/whats-in-a-search-engines-name

Simon, M. (2019, March 11). The genderless digital voice the world needs right now. *Wired.* https://www.wired.com/story/the-genderless-digital-voice-the-world-needs-right-now/

Sin, S. & Munteanu, C. (2019). An information behaviour-based approach to virtual doctor design. *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Service*s. ACM Digital Library .DOI:/10.1145/3338286.3344391

Smith, A. (2018, June 28). Public attitudes toward technology companies. *Pew Research Center.* https://perma.cc/KSN6-6FRW

Smith, W. G. (2008). Does Gender Influence Online Survey Participation? A Record-Linkage Analysis of University Faculty Online Survey Response Behavior. *Online Submission*. https://files.eric.ed.gov/fulltext/ED501717.pdf

SRI International (n.d.). 2007 Siri [website].Retrieved December 15, 2021, from https://www.sri.com/hoi/siri/

Stahl, B. C., Timmermans, J., & Mittelstadt, B. D. (2016). The ethics of computing: A survey of the computing-oriented literature. *ACM Computing Surveys 48* (4). DOI: 10.1145/2871196

Statistics Canada (2016a). Age and sex highlight tables, 2016 Census. https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/hlt-fst/as/Table.cfm?Lang=E&T=11

Statistics Canada (2016b). Census profile, 2016 census Alberta [province] and Canada [country]. https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/prof/details/Page.cfm?Lang=E&Geo1=PR&Code1=48&Geo2=&Code2=&Data=Count&SearchType=Begins&SearchPR=01&B1=All

Statt, N. (2018, May 10). Google now says controversial AI voice calling system will identify itself to humans. *The Verge.* https://www.theverge.com/2018/5/10/17342414/google-duplex-ai-assistant-voice-calling-identify-itself-update

Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online manipulation: Hidden influences in a digital world. *The Georgetown Law Technology Review 4*(1). https://link.gale.com/apps/doc/A613511204/LT?u=edmo69826&sid=LT&xid=4d782d75

Sydell, L (2018, July 9) The push for a gender-neutral Siri. *NPR* https://www.npr.org/2018/07/09/627266501/the-push-for-a-gender-neutral-siri

Tassiello, V., Tillotson, J.S., Rome, A.S. (2021). "Alexa, order me a pizza!": The mediating role of psychological power in the consumer–voice assistant interaction. *Psychology & Marketing, 38*(7). DOI:https://doi.org/10.1002/mar.21488

Tene, O., & Polonetsky, J. (2014). Theory of creepy: Technology, privacy and shifting social norms. *Yale Journal of Law and Technology, 16,* 59-102. https://heinonline.org/HOL/P?h=hein.journals/yjolt16&i=60

Tigard, D. W., Conradie, N. H. & Nagel, S. K. (2020). Socially responsive technologies: toward a co-developmental path. *AI & Society, 35*, 885–89. DOI:10.1007/s00146-020-00982-4

Université de Montréal (2018). Montréal declaration for a responsible development of artificial intelligence. https://www.montrealdeclaration-responsibleai.com/the-declaration

Vigliarolo, B. (2020, June 10). Alexa skills: Cheat sheet. *TechRepublic.* https://www.techrepublic.com/article/alexa-skills-cheat-sheet/

Vimalkumar M., Sharma S. K., Singh J. B. & Dwivedi Y. K. (2021). 'Okay Google, what about my privacy?': User's privacy perceptions and acceptance of voice based digital assistants. *Computers in Human Behavior, 120.* DOI: 10.1016/j.chb.2021.106763.

Vincent, J. (2018, May 9). Google's AI sounds like a human on the phone- should we be worried. *The Verge.* https://www.theverge.com/2018/5/9/17334658/google-ai-phone-call-assistant-duplex-ethical-social-implications

Vixen Labs, Open Voice Network, & Delineate (2021, June 30). US voice assistant users, by age and brand. *eMarketer.* https://www.emarketer.com/chart/249362/us-voice-assistant-users-by-age-brand-may-2021-of-respondents-each-group

Vehovar, V., Toepoel, V. & Steinmetz, S. (2016). Non-probability sampling. In C. Wolf, D. Joye, T. W. Smith, & Y. Fu (Eds.), *The SAGE Handbook of Survey Methodology* (pp. 329-345). SAGE Publications Ltd. DOI: 10.4135/9781473957893

Wakefield, R. (2013). The influence of user affect in online information disclosure. *The Journal of Strategic Information Systems, 22*(2), 157–74. DOI: 10.1016/j.jsis.2013.01.003

Waytz, A., Klein, N., & Epley, N. (2013). Imagining other minds: Anthropomorphism is hair-triggered but not hare-brained. In M. Taylor (Ed.) *The Oxford Handbook of the Development of Imagination.* DOI: 10.1093/oxfordhb/9780195395761.013.0018

Whitby, B. (2012). Do you want a robot lover? The ethics of caring technologies. In P. Lin, K. Abney, & G.A. Bekey (Eds.), *Robot ethics: The ethical and social implications of robotics* (1st ed., pp. 233-248). Cambridge, Mass.: MIT Press. https://ebookcentral.proquest.com/lib/ualberta/reader.action?docID=3339351&ppg=216

Williams, H. T. P., McMurray, J. R., Kurz, T., & Lambert, F. H (2015). Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Global Environmental Change*, *32*, 126–138. DOI: 10.1016/j.gloenvcha.2015.03.006

Woods, H. S. (2018). Asking more of Siri and Alexa: Feminine persona in service of surveillance. *Critical Studies in Media Communication, 35*(4), 334-349. DOI: 10.1080/15295036.2018.1488082

Xie, W., Fowler-Dawson, A., & Tvauri, A. (2019). Revealing the relationship between rational fatalism and the online privacy paradox. *Behaviour & Information Technology*, *38*(7), 742–759. DOI:10.1080/0144929X.2018.1552717

Yang, H., & Lee, H. (2019). Understanding User Behavior of Virtual Personal Assistant Devices. *Information Systems and E-Business Management*, *17*(1), 65–87. DOI: 10.1007/s10257-018-0375-1

Zhang, B., and Dafoe, A. (2019). Artificial intelligence: American attitudes and trends. Center for the Governance of AI, Future of Humanity Institute, University of Oxford. https://governanceai.github.io/US-Public-Opinion-Report-Jan-2019/us_public_opinion_report_jan_2019.pdf

Zuboff, S. (2015). Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology, 30*(1), 75-89. DOI:10.1057/jit.2015.5

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power.* PublicAffairs.

Appendix A: Survey

Thank you for your interest in participating in this survey. This research project is being conducted by Julia Guy, a Masters in Digital Humanities and Library and Information Studies (MA/MLIS) candidate at the University of Alberta. This research aims to investigate public opinion about Virtual Assistants (VAs) such as Apple's Siri and Amazon's Alexa. VAs are accessible most commonly through phones, smart speakers, and in the car. They are designed to respond to people's voices in a conversational way. You do not need to be a regular user of these technologies to participate but you should have a basic understanding of what they are and how they are used.

The aim of this research is to collect a variety of different opinions on Virtual Assistants and in particular how these technologies portray human-like qualities. The analysis process will involve identifying any trends in public opinion that may appear in the data and interpreting those results. Responses to the survey will be used as a jumping-off point for considering the ethical implications of these technologies and how that relates to public opinion.

Anyone 18 years old or older may participate. The survey will take approximately 10-20 minutes to complete. This survey is anonymous and no identifying information will be collected. Participation is voluntary and nobody will be able to tell whether you have completed it or not. If you change your mind about wanting to participate at any point while you are completing the survey simply close the browser window without clicking on 'submit.' Once responses are submitted they cannot be edited or removed from the data because they will not be identifiable. Please note that by completing this survey you are providing consent for your responses to be recorded and analyzed by the researcher. The collection of survey responses will close on August 12th 2021, so if you are interested in participating please do so before then.

The plan for this study has been reviewed and approved by the Research Ethics Board at the University of Alberta (application ID: Pro00110059). If you have any questions or concerns regarding this research or your rights as a participant you may contact the University of Alberta's Research Ethics Office at 780-492-2615. This office has no affiliation with the researcher. If you would be interested in seeing the results of this survey you can contact Julia Guy after March 2022. You can also contact the researcher at any time for further information or clarification.

Contact information:

Researcher: Julia Guy
jguy@ualberta.ca
(587) 589-8706

Supervisor: Tami Oliphant
toliphan@ualberta.ca
(780) 492-2033

Supervisor: Geoffrey Rockwell
geoffrey.rockwell@ualberta.ca
(780) 492-6436

Section One: Use of Social AI [Note: Section titles were not visible to respondents]

What is your age:

- ❑ 18-29
- ❑ 30-39
- ❑ 40-49
- ❑ 50-59
- ❑ 60-69
- ❑ 70+

What is your gender (this information will not be used to identify participants):

- ❑ Agender (does not relate to the idea or experience of gender)
- ❑ Bigender (both male and female at varying points)
- ❑ Female
- ❑ Male
- ❑ Non-Binary (does not align with the man-woman binary)
- ❑ Transgender Man
- ❑ Transgender Woman
- ❑ Two-spirit
- ❑ Would rather not say
- ❑ Other: _____

Do you currently live in Canada?

- ❑ Yes
- ❑ No

Have you used Virtual Assistants such as Alexa, Siri, Cortana, or Google Assistant/ Home? Select all of the technologies you have interacted with:

- ❑ Amazon's Alexa
- ❑ Apple's Siri
- ❑ Google Assistant

- ❏ Microsoft Cortana
- ❏ None
- ❏ Not sure
- ❏ Other:_____

How often do you interact with a Virtual Assistant?

- ❏ Once a month or less
- ❏ Roughly every two weeks
- ❏ Roughly once a week
- ❏ Several times a week
- ❏ Everyday
- ❏ Multiple times a day
- ❏ I used to use them but I don't anymore
- ❏ I have never used one
- ❏ Other: _____

Where do you use Virtual Assistants? Select all that apply.

- ❏ At home
- ❏ At work
- ❏ In the car
- ❏ Outside the home
- ❏ Traveling
- ❏ No place in particular
- ❏ I used to use them but I don't anymore
- ❏ I have never used one
- ❏ Other :_____

What do you use a Virtual Assistant for? Select all that apply.

- ❏ Checking your schedule
- ❏ Fun (asking funny questions etc.)
- ❏ Getting directions

- ❏ Making lists
- ❏ Making phone calls
- ❏ Playing audio or video content
- ❏ Playing games
- ❏ Searching the internet
- ❏ Seeking recommendations
- ❏ Sending text messages or emails
- ❏ Smart home activities such as adjusting temperature, etc.
- ❏ None, because I do not use Virtual Assistants
- ❏ Other:_____

Using Virtual Assistant technology is: Select all that apply.

- ❏ Amusing
- ❏ Annoying
- ❏ Boring
- ❏ Creepy
- ❏ Confusing
- ❏ Convenient
- ❏ Difficult
- ❏ Draining
- ❏ Easy
- ❏ Empowering
- ❏ Fine (no strong feelings about it)
- ❏ Frustrating
- ❏ Fun
- ❏ Helpful
- ❏ Satisfying
- ❏ Nerve-wracking
- ❏ Other:_____

On a scale from 1 to 10, how aware would you say you are about how Virtual Assistants store and use personal data?

Do you use location services on our devices?

- ❑ Yes
- ❑ No
- ❑ Not sure

Section Two: Concerns

Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact people?

- ❑ Yes
- ❑ No
- ❑ Somewhat

Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact society?

- ❑ Yes
- ❑ No
- ❑ Somewhat

Do you trust that technology companies that make Virtual Assistants are making sure their technologies do not negatively impact the environment?

- ❑ Yes
- ❑ No

- ❑ Somewhat

What is your overall level of concern about surveillance and data collection by tech companies (including social media companies, search engines, and other platforms)?

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

What is your overall level of concern about surveillance and data collection specifically by human-like AIs such as Virtual Assistants?

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

What specific ethical concerns, if any, do you have about Virtual Assistants? Rate the level of concern you have for each of the following:

Being stereotyped/profiled:

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Consumerism/buying things you don't need

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Corporate surveillance

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Discrimination

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Distribution of audio recordings

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Gender stereotypes

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned

- ❑ Very concerned
- ❑ Extremely concerned

Government surveillance

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Increased political polarization

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Inequality in access and usability of these technologies

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Lack of anonymity (having your data identifiable to you)

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Lack of transparency by companies

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Loss of human connection

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Manipulation by companies

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Political manipulation

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Technology impersonating humans

- ❑ Not at all concerned

- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Risk of VA being hacked

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Risk of privacy breach

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Technology dependency

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Threats to autonomy (the ability to make your own decisions)

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned

- ❑ Extremely concerned

VAs taking human jobs

- ❑ Not at all concerned
- ❑ A little concerned
- ❑ Moderately concerned
- ❑ Very concerned
- ❑ Extremely concerned

Are there any other concerns you have that are not listed above?

_____

What do you think is the most positive thing about Virtual Assistant technology?

_____

What do you hope Virtual Assistants will be like in the future?

_____

Section Three: Anthropomorphization

What information would you rather get from a Virtual Assistant rather than a real person? Check all that apply.

- ❑ Basic medical or health related information
- ❑ Complex medical or health related information
- ❑ Directions
- ❑ How-to information
- ❑ Leisure and restaurant recommendations
- ❑ Media and entertainment recommendations
- ❑ News
- ❑ Schedule details
- ❑ Shopping assistance

- ❏ Simple facts (e.g. 'What is the capital of Manitoba?)
- ❏ Translation services
- ❏ None
- ❏ Other:_____

How human-like should Virtual Assistants be?

- ❏ Not human-like at all
- ❏ Only human-like enough that you can communicate with it verbally but not more
- ❏ Human-like enough that talking to it it feels like you are having conversation with a real person
- ❏ Human-like enough that you would think of that Virtual Assistant as a real person
- ❏ Human-like enough that you consider it a close friend
- ❏ Other:

What human-like traits should Virtual Assistants have? Select all that apply

- ❏ Assertive
- ❏ Clear
- ❏ Clever
- ❏ Considerate
- ❏ Discreet
- ❏ Eager to please
- ❏ Energetic
- ❏ Helpful
- ❏ Humourous
- ❏ Fair
- ❏ Firm
- ❏ Friendly
- ❏ Kind
- ❏ Motivational
- ❏ Nurturing
- ❏ Opinionated

- ❏ Polite
- ❏ Positive
- ❏ Reliable
- ❏ Reserved
- ❏ Sassy
- ❏ Smart
- ❏ Submissive
- ❏ Thoughtful
- ❏ Tells it like it is
- ❏ Trustworthy
- ❏ Unique
- ❏ No human-like traits
- ❏ Other:_____

What gender should Virtual Assistants portray?

- ❏ Agender (does not relate to the idea or experience of gender)
- ❏ Bigender (both male and female at varying points)
- ❏ Female
- ❏ Male
- ❏ Non-Binary (does not align with the man-woman binary)
- ❏ Two-spirit
- ❏ Does not matter
- ❏ Other: _____

What age should Virtual Assistants portray?

- ❏ Under 20
- ❏ 21-30
- ❏ 31-40
- ❏ 41-50
- ❏ 51-60
- ❏ 61-70

- ❑ 71+
- ❑ Does not matter
- ❑ Other: _____

Section Four: Data collection

Who should be responsible for making sure Social AIs like Siri and Alexa are developed in a way that does not negatively affect people or society? Select all that apply.

- ❑ Governments
- ❑ Individuals
- ❑ International organizations
- ❑ The companies themselves
- ❑ Universities and researchers
- ❑ All of the above
- ❑ Nobody
- ❑ Other:_____

Should a Virtual Assistant have to identify itself as AI?

- ❑ Yes
- ❑ No
- ❑ It depends

If you selected "It depends" for the last question could you elaborate?

Should the data collected by a Virtual Assistant be used to help advertisers profile users for targeted ads?

- ❑ Yes
- ❑ No
- ❑ It depends

If you selected "It depends" for the last question could you elaborate?

Should Virtual Assistants be able to recommend products or businesses to you?

- ❑ Yes
- ❑ No
- ❑ It depends

If you selected "It depends" for the last question could you elaborate?

Should Virtual Assistants be able to share your location with third-parties?

- ❑ Yes
- ❑ No
- ❑ It depends

If you selected "It depends" for the last question could you elaborate?

Should Virtual Assistants be able to share recordings of users with company employees in order to improve services?

- ❑ Yes
- ❑ No
- ❑ It depends

If you selected "It depends" for the last question could you elaborate?

Should the data collected by a Virtual Assistant be used to help curate people's social media feed?

- ❏ Yes
- ❏ No
- ❏ It depends

If you selected "It depends" for the last question could you elaborate?

When should law enforcement or the government be able to access personal information gathered by a Virtual Assistant? Select all that apply.

- ❏ Never
- ❏ When they deem it is absolutely necessary
- ❏ When children are at risk
- ❏ When victims of crimes request an investigation
- ❏ Anytime
- ❏ Other:_____

When should Virtual Assistants be able to collect and save data and recordings:

- ❏ Only when people have activated it
- ❏ Anytime
- ❏ Not all the time
- ❏ Never
- ❏ Other:

On a scale of 1 to 10 how much information do you want about how Virtual Assistants use, collection, and distribute personal data?

	1	2	3	4	5	6	7	8	9	10	
None	○	○	○	○	○	○	○	○	○	○	Absolutely everything

How do you want this information about how your data is managed communicated to you? Select all that apply

- ❏ I don't want this information
- ❏ In a short video
- ❏ In a standard Terms and Conditions agreement
- ❏ In the voice of the Virtual Assistant itself
- ❏ In simple language
- ❏ I want to be notified every time my data is shared
- ❏ I want to talk to a real person about it
- ❏ No preference
- ❏ Over text message
- ❏ Written somewhere clearly on the technology's website
- ❏ Other:_____

Do you have any other thoughts about Virtual Assistants that you would like to mention?

_____

Appendix B: Percentages of Age Groups in Alberta (Statistics Canada, 2016b)

The chart below shows the percentage of the population of Alberta in 2016 (not including children under 10) for each age group outlined as an option for respondents to select in the survey. This research aimed to have as similar a distribution of respondents ages as possible to the information outlined in the chart (E.g. have roughly 20% of respondents be ages 30-39).

Table B1

Percentage of the 2016 Alberta Population for Each Age Group Option Provided to Respondents

Total Population of Alberta in 2016: 4,067,175
Population not including those ages 0-18: 3,144,381

Age	**Population**	**Approximate % of Population (not including those ages 0-18)**
0-18 (will not be surveyed for this research)	922, 794	N/A
18- 29	669 162 *	21%
30-39	638,680	20%
40-49	550,235	18%
50-59	559,850	18%
60-69	399,910	13%
70+	326,540	10%

** Note: The number of people 18-29 (669 162) was determined by taking 20% of the population of 10-19 year olds (96,391) and adding that to the population of 20-29 year olds (572,770). This approximation of the population of 18-29 year olds was used because only consenting adults will be asked to complete the survey.*

Appendix C: Survey Distribution Plan

The following tables outline how participants were recruited through a combination of listservs, social media promotion, and by word-of-mouth or snowball sampling. Those asked to help distribute the survey were also encouraged to share it with people who might have different opinions than their own.

Social Media

Platform	Pages	Reasoning
Facebook	● Personal Account (created for this research) ● Volunteer Alberta	● Because I do not have followers and am just joining Facebook for the duration of this research, this platform will be primarily used to distribute the survey to those who would likely not use the other platforms (likely those over 50)
Reddit	● SampleSize ● University of Alberta ● Technology	● SampleSize (161k subscribers (not Alberta specific) ● University 25.8K subscribers ● Technology is a subreddit dedicated to the news and discussions about the creation and use of technology and its surrounding issues and may therefore have people interested in completing the survey
Twitter	● Personal account	
Instagram	● Personal account	

Personal Contacts

Person/People	Reasoning

Family	There is a good chance that family members will respond and share the survey. Family members belong to a range of different communities and have professional and personal contacts of various ages and professions including warehouse workers, the online gaming community, teachers, artists, bankers, people working in the non-profit sector etc.
Friends	There is a good chance that friends will respond and share the survey. Friends of different ages and from different contexts will be asked to try to increase variety in responses. They will also be encouraged to share it with relatives of different ages and on social media if they choose.
DH Cohort	Although there is a greater chance that my fellow students may have views that align with mine and each other, our shared interest in the topic of this research will probably add to the likelihood that they will participate and share the survey. The research will be explained in more detail with this group because they are likely to understand the sample I am trying to achieve and help share the survey with people and groups that I may not have access to on my own.
SLIS Cohort	See above.
Restaurant Co-workers	The survey will be shared in the staff chat at the restaurant I work at in Edmonton. Staff are mostly between 18 and 40 years old with a few exceptions, however their professional goals, interests and experience with technology have a lot of variety.
Thesis Committee	If they are comfortable doing so, my thesis committee members will be asked to share the survey with their contacts.
BFA Acting Class 2012	My undergrad cohort will be asked to share the survey primarily because they are likely to do so, many have many social media followers, and several have moved into different professions such as law and social work and may have different contacts to share it with.
Theatre community contacts	Theatre community contacts such as cast-members and acquaintances will be asked to complete and share the survey. This is partly due to the fact that most of my friends that are over 50 years old are people that I have met working in the professional theatre. Many of whom I think would be excited to complete and share the survey with their friends. Theatre professionals have varying experience with technology and research and generally exist outside of the academic sphere which is where many of my other participants will be from.

Listserves

Listserve	**Reasoning**

Jerome	JEROME-L is hosted by the School of Library & Information Studies and was created to foster communication among librarians, information professionals, and friends of libraries in Alberta. Recruiting participants would likely be approved for distribution through this listserv and hopefully people would be willing to participate.
Digital Humanities	HUCO-grad-students listserve will allow me to reach my DH cohort. I will also post it on our student server on Discord.

Appendix D: Concerns About VAs

To identify which issues participants found the most concerning out of the 20 options provided in the survey, the percentage of respondents who were "Very concerned" or "Extremely concerned" for each issue was calculated.

Table D1

Full List of VA Concerns, Rated in Order of Percentage of Respondents Who Were "Very" or "Extremely" About Each Issue

Specific Concerns	Percentage of respondents either "Very concerned" or "Extremely Concerned"
Lack of Transparency by Companies	67.8%
Manipulation by Companies	63.3%
Increased Political Polarization	60%
Corporate surveillance	55.6%
Political Manipulation	55.6%
Lack of anonymity (having your data identifiable to you)	54.4%
Risk of Privacy Breach	46.6%
Consumerism/buying things you don't need	41.1%
Inequality in access and usability of these technologies	41.1%
Risk of VA being hacked	37.8%
Technology impersonating humans	35.6%
Decrimination	35.6%
Loss of human connection	34.4%
Distribution of audio recordings	33.3%

Gender stereotypes	33.3%
Government surveillance	32.2%
Being stereotyped/profiled	32.2%
Threats to autonomy (the ability to make your own decisions)	28.9%
Technology dependency	27.8%
VAs taking human jobs	23.3%