From Virtual Assistants to AI: Data Privacy Issues in the Digital Age

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

William Harrison Tan

Spring 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Joshua Earle, Department of Engineering and Society

From Virtual Assistants to AI: Data Privacy Issues in the Digital Age

Introduction

In 2017, Equifax, of the three big credit bureaus in the United States, suffered a data breach that compromised the personal information of over 160 million people. The hack compromised critical information, such as names, birth dates, social security numbers, home addresses, and license numbers, giving rise to the possibility of their being used for illicit activities such as identity theft. It was, as then Attorney General Barr described it, one of the ’largest data breaches in recent history.’ (“Data from Equifax”, 2020) Though it has been slightly lost in the noise, being crowded out by other data breaches in more recent years, it is arguably one of the most significant incidents when it comes to the breadth and depth of the information that was compromised. Companies in the present day, experts say, collect far too much information, much of it sensitive, and make that information vulnerable to falling into the wrong hands (“Data from Equifax”, 2020). This instance was a stark reminder as to what information should be entrusted to companies such as Equifax, as well as what information people are giving up without their knowledge.

Currently, the data being collected from people as well as the privacy of such data have both become increasingly important issues. Not limited to what users fill out in forms or in search bars, today’s technology records trends, habits, and other abstract - yet deeply personal - aspects of their users. With ever-improving metrics that can more accurately summarize what was once thought to be too-complex processes limited to the domain of humans, data has become more valuable than ever. The latest developments in virtual assistants, as well as artificial intelligence in general, have emphasized the tremendous uncertainty and opacity around how users’ data is used. Consequently, the most money to be made in tech revolves around data and its endless, constant harvest.

The rapid interconnectedness of our world has led to the widespread sharing of personal data, often without individuals fully understanding the terms of service they agree to. Users are unaware and possibly uninformed about how they can risk their data privacy while participating, especially when interacting with virtualized assistants and, more recently, artificial intelligence. In other cases, users are well aware of the data they may be giving away, but are largely resigned to doing so and feel powerless to change such a situation. In this paper, I will explore the unique vulnerabilities people get when interacting with virtual assistants and artificial intelligence. I will then further explore the possible social contract relationship between such applications, their developers, and the users. Afterward, I will examine the current regulatory and cultural landscape for such things and propose possible ways forward when it comes to regulation.

The Social Contract

As Fred D’Agostino and coauthors put it, Hobbs held the structure of the social contract as follows: people give up certain freedoms to society in exchange for common interests such as protection. This is often referred to when discussing the forms, functions, and responsibilities of governments, most of which exist due to the consent of the governed. (D’Agostino et al., 2017). However, by considering the social contract as a framework through which one party gives up something in exchange for something else, we can apply it in a slightly different light here: people give up certain information about themselves in exchange for the use of a virtual assistant or artificial intelligence. Most people are conscious that they are giving up some information, as they know that the agreements and contracts they click through to access a website or use an application surely mean that they give up something in exchange.

Some types of information that will be given up are obvious: anyone using a virtual assistant or AI knows that they may be giving up clips of their voice, the text they write, and the interests they show it to allow them to give advice or make suggestions. However, it is rarely actually that simple. Kaveh Waddell writes that there are also more subtle types of information that people may be giving up. This includes their location data, browsing history, and even the patterns in which they interact with their devices, all of which are captured with the purpose of improving software quality and efficiency (Waddell, 2016). This information can reveal a lot about a person, including their daily routine, their interests, and their personality traits, and is well beyond most users' understanding. In many cases, this information is collected and analyzed without the user even realizing it, as it happens in the background of the devices and services they use, as Geoffrey Fowler analyzed in the cases of Alexa and Siri. Fowler also found that these virtual assistants would often phone home to their respective companies, transmitting recorded information in the middle of the night. The users were completely unaware that this happened (Fowler, 2019a).

One of the most important aspects of the social contract is that of informed consent, the consent of the person upon their being informed of what they would be giving up in exchange for a service. Most users, as established by Brooke Auxier and coauthors, only have a very vague or abstract view of what they actually do sign away (Auxier et al., 2019). Most users, when agreeing to use these kinds of software and applications, know and understand that they hand over obvious information: a note, a voice clip, or a spot of camera footage. However, they are unaware of the extent to which their personal data is being collected, analyzed, and potentially shared with third parties. This lack of transparency and understanding is incompatible with informed consent, which forms the foundation on which the social contract is made.

Along with informed consent, the right to exit is also a crucial but less-discussed element of a functional social contract. Since a social contract implies a freedom of association, or to consciously choose which system is right for a user, it also implies a freedom to disassociate or to leave in order to find a better arrangement. As Élise Rouméas puts it, “[it is] the right to end an association without excessive and undue cost.” (Rouméas, 2023) However, in this present era, every arrangement a user can enter into will have similar practices when it comes to user data. The wide adoption of these kinds of data-collecting processes by the majority of the tech industry, as well as the lack of transparency about what is collected and its uses, have led to a degree of pessimism and resignation among users since the right to exit does not effectively exist – not unless a user decides to stop being a user. This extreme course of action is also exceptionally costly, as a user would have to give up the benefits of using highly convenient software and even hardware to finally be free from this agreement. Modern discussion of the social contract also seeks to ensure that people under it should always be free to leave and are not coerced into remaining bound by it (Rouméas, 2023), and this is a glaring concern when using the social contract as a framework to examine user-company relationships.

The transactional nature of the social contract allows for an effective examination of today’s usage of virtual assistants and artificial intelligence. In this context, both of the two most important parts of the social contract are in jeopardy. Users are not informed nor can they pursue more suitable alternatives without having to give up a large amount of technology usage, suffering an undue cost in the process – and therefore have no real right to exit. The two case studies below both illustrate and emphasize the issues with the current social contract, as well as the plight of the software’s users in being uninformed and unable to leave.

Case Study: Virtual Assistants

Before artificial intelligence and its endless appetite for data became rooted in popular culture, it was preceded by virtual assistants. These programs and applications worked with natural language - language that everyday people use. With a single command, as Erick Schonfeld puts it, a user could see the weather forecast, check stock prices, make reminders, and set appointments on a calendar (Schonfeld, 2010). When Siri was first released in 2010 and put both virtual assistants and artificial intelligence in the wider public spotlight, it opened the door for far more improvement and far more visibility for the issue of data privacy.

Though virtual assistants first had to be prompted by the press of a button or the typing of a command, later ones could be spoken to. With a single spoken prompt, such as a "hey" or a "hello," a virtual assistant could be brought online and ready to do what the user would ask of it. One major concern is that virtual assistants may be listening to users' conversations without their knowledge or consent. According to Kieren McCarthy, reports emerged in 2018 that Alexa had recorded a private conversation and sent it to someone who was entirely unrelated . When prompted through the European General Data Protection Regulation (GDPR), Amazon provided an Alexa user with what was apparently recordings attributed to him – only to send him the recordings of a complete stranger (McCarthy, 2018). This incident raised questions about whether virtual assistants are always listening, and whether they are able to distinguish between intentional commands and casual conversation. It also raised concerns over how the data collected are handled and used by the companies that developed the virtual assistants.

Another issue is the usage of users' personal data collected by virtual assistants, including how it is used and who it is used by. In 2019, it was revealed that Amazon had been sharing users' Alexa voice recordings with contractors to improve the accuracy of responses (Valinsky, 2019). While Amazon claimed that these contractors were subject to strict privacy policies, the incident raised concerns about how much control users have over their personal data. Similarly, according to Hern, Apple also employed third-party contractors in order to evaluate voice recording data, with contractors ending up with data that was supposed to be confidential, such as doctor-patient conversations, or was extremely personal, such as couples having sex (Hern, 2019). Similar issues cropped up for Google, as well (Clauser, 2019). In any of the controversies, such usage of data was not apparent in user agreements and could not be opted out from, emphasizing the problems plaguing the social contract between the developers and the users.

Furthermore, virtual assistants are often integrated with other devices, such as smart speakers and home automation systems, which increases the amount of personal data that can be collected (Fowler, 2019b). This data can be used to build detailed profiles of users, which could be exploited for targeted advertising or even identity theft (Levin, 2017). For instance, Amazon's Alexa system constantly listens to ambient noise and conversation, sending the data back to Amazon to recommend particular items to the user. As documented in an investigation, an Alexa would listen and record audio every time it thought it heard its name, sending data that it was absolutely not supposed to record (Lerman, 2021). Tom Bolton and coauthors indicate that that most users are unaware that their devices, when integrated into a voice assistant system, would be constantly recording and transmitting in order to collect information about them. They also indicate that if the company suffers a breach, such information can be used in order to mimic voices and impersonate a particular user when used in conjunction with other compromised details (Bolton et al., 2021). Voice and audio recordings can contain far more insight than what it looks like on the surface, according to Jacob Leon Kröger and coauthors, who show that things such as "biometric identity, personality, physical traits, geographical origin, emotions, level of intoxication and sleepiness, age, gender, and health condition." (Kröger et al., 2020) Even when apparently stripped of personal identifiers for regulatory purposes, as Amazon, Google, and Apple have publically stated (Clauser, 2019), voice data can still be effectively used to gain information about its source even when ostensibly unlabelled.

When considering the issue of virtual assistants through the social contract framework, the issues of privacy are immediately apparent. Virtual assistants have access to a vast amount of extremely personal data and information that users share with them often without awareness that such information and data are even being recorded. As such, the use of virtual assistants can raise serious concerns about the privacy of individuals, and how their personal data is being collected, stored, and used. To address these concerns, it is essential to consider the social contract framework and virtual assistant providers' obligations to their users (Agrawal et al., 2018). This includes ensuring transparency in data collection and usage, obtaining explicit consent from users, and providing mechanisms for users to control and manage their personal data (Cho et al., 2020). Failure to uphold these obligations can erode the trust between users and virtual assistant providers, and continue to undermine the social contract that underpins their relationship.

Case Study: The Era of Artificial Intelligence

Artificial intelligence, or AI, took the technological world by storm in the mid-to-late 2010s and is currently experiencing a meteoric rise in usage and popularity (Tiku et al., 2023). The ability of machines, algorithms, and programs to learn and make intelligent decisions revolutionized industries ranging from healthcare to finance to transportation, and AI became more accessible to businesses of all sizes with the proliferation of big data, cloud computing, and advanced algorithms. This resulted in the development of new products and services that have transformed the way people live and work (Abril, 2023). However, the widespread use of AI has also opened up new data privacy concerns, particularly that of what actual data most AIs and AI companies use to improve and develop, respectively.

While popular conceptions of AI previously revolved around programs and algorithms that could rival humans in intellect and understanding, AI already permeates an incredible number of technologies even in a relatively basic form. With advancements in natural language processing, computer vision, and machine learning, AI has become increasingly sophisticated and powerful. The potential for AI to transform industries and solve complex problems is immense. However, the power of AI is also evident in its ability to process vast amounts of data, which has raised serious concerns regarding data privacy. As AI continues to evolve and become more integrated into our lives, it will be essential to ensure that data privacy is protected and that users are fully informed about how their data is being used. For instance, according to Cameron F. Kerry, companies should be held to account via measures for transparency, explainability, auditing, and risk assessments (Kerry, 2020). This way, users can consider their data to be reliably safeguarded and they can also be informed in the ways above.

AI requires an ever-increasing amount of data to develop, which leads to more data being harvested from users. Recently, as AI development began to take off, the race to create better AI has seen developers become more reckless and risky, exchanging safety for speed (Kerry, 2020). Better AI, however, requires more data, which in turn requires more data harvesting from large portions of the internet (Gravrock, 2022). This has raised concerns about the ethical implications of using personal data in AI development and the potential consequences for user privacy. The rise of automated data collection tools, such as web scraping bots, has added to these concerns, as they can be used to collect data on a massive scale without users' knowledge or consent. After all, the usage of personal data and the details that can be derived from it for training an AI of unknown purpose would be a gross violation of the social contract.

Entire industries have been created to simply collect data from users, operating under agreements and procedures a user has little to no idea about. The collection and use of personal data by these companies have raised concerns about data privacy and the potential for misuse. As AI continues to advance, it will be essential to establish clear regulations and guidelines to ensure that personal data is collected, used, and stored in an ethical and responsible manner. This will require collaboration between governments, industry, and the public to develop policies that balance the benefits of AI with the protection of individual privacy.

While AI was still largely in its adolescence in the middle of the 2010s, the Cambridge Analytica scandal is probably one of the best-known examples of ethically questionable usage of both AI and data collection. The use of AI in the Cambridge Analytica scandal involved the development of algorithms to analyze and predict the behavior of Facebook users based on their personal data, and the organization worked to harvest such data, discovering trends, strengths, and vulnerabilities of groups of people. It would then use such information to influence political campaigns in various nations. The biggest example was its involvement in the 2016 United States presidential election, where it delivered targeted advertisements to specific groups, likely influencing the way they voted (Chang, 2018). The scandal raised questions about the role of AI in political campaigns and the need for transparency in data collection and use (Lapowsky, 2019). It also highlighted the potential risks of data misuse by third-party companies and the need for stronger data privacy regulations.

Later, in 2019, a facial recognition start-up called Clearview AI was found to be collecting and selling personal data without user consent (Heikkilä, 2022). Clearview AI had created a massive database of people's faces by scraping images from social media sites, including Facebook and Twitter (Hill, 2020). It would then use this data to develop facial recognition AI that could identify people based on their faces. However, the company's data collection practices and the use of facial recognition technology raised serious ethical concerns about privacy and surveillance. The pictures that were used were obtained without consent, drawing attention to the usage of personal data for potentially nefarious purposes.

The controversies around Cambridge Analytica and Clearview AI demonstrate the need for additional transparency and accountability in both AI development and data collection. The fact that companies are able to harvest personal data and put it to any kind of usage without users' knowledge or consent is a clear violation of the social contract between users and companies. As part of the social contract, users trust companies to protect their personal data and use it only for legitimate purposes. In the current state of the art, this has not been happening, with regulations and policies governing such an arrangement only recently being proposed. This is especially true in the instance of ChatGPT, which shares similar controversies when it comes to data privacy and users' unawareness of it (Johnson-Gomez, 2023). In fact, even ChatGPT suffered through both growing pains, where users were able to see the conversations of other users, and data breaches, where personal data was once again spilled out into the open web (Abrams, 2023). With AI starting to take an outsize role in society, it is absolutely critical for users to have a say in how their data is used.

Discussion

The case studies of virtual assistants and AI above highlight the urgent need for improved data privacy measures. The rapid pace of technological innovation has made the usage of users' data ever more important to acquire, which drives companies and developers to implement ever more complex measures to harvest it. This has created a situation where companies and developers must balance the benefits of AI against users' privacy concerns. The current state of data privacy regulations and practices has been unable to keep pace with the speed of technological innovation, leading to situations where companies have taken liberties with users' data and put it to arguably ethically dubious purposes. This has violated the social contract between users and companies, where users presume their data and information would be used in predictable and understandable ways. Therefore, as AI continues to evolve, it's crucial that companies and regulators prioritize data privacy measures to ensure that users' rights and interests are protected.

The current relationship between users and developers is unbalanced, emphasizing the importance of informed consent and the right to exit in creating a functional social contract. To address this for instance, companies must prioritize user privacy and provide transparent information about data collection practices, including how data is being used and shared (Medairy). Additionally, users must have the ability to opt-out or exit from data collection at any time, without negative consequences. This right to leave, or right to delete, would create a more equitable relationship between users and companies, allowing users to have greater control over their personal data and ensuring that companies respect users' privacy rights.

We can propose several possible measures to allow users to become more informed or to have an easier ability to opt-out or exit, which are the two concepts critical to a functional social contract (King and MacKinnon, 2022). One approach could be to increase user awareness about data collection and usage, providing clear and concise explanations of how personal data is collected and used by companies. Another option could be to provide users with more control over their personal data through the use of privacy-enhancing technologies, such as encryption or decentralized data storage. In cases where privacy breaches occur, third-party intervention could be considered as a solution, allowing users to seek assistance and redress from independent organizations that specialize in privacy protection. These measures would help create a more balanced relationship between users and companies, allowing users to have greater control over their personal data and ensuring that companies respect users' privacy rights.

In conclusion, the rapid pace of technological innovation has created an urgent need for the mending of the social contract and improved data privacy measures. As AI development continues to accelerate, companies and regulators must prioritize user privacy and establish clear guidelines to ensure that users' rights are protected. This requires a shift in mindset, where companies view data privacy as a fundamental right, rather than an afterthought. Users must be empowered with informed consent and the ability to opt-out or exit from data collection, while companies must be transparent about their data collection practices and accountable for any misuse of personal data. Achieving a functional social contract between users and companies requires a collaborative effort from all stakeholders, including regulators, industry leaders, and users themselves. By working together to prioritize data privacy, society can ensure that the benefits of AI are balanced with the protection of users' privacy rights. The future of AI depends on it.

References

Abrams, L. (2023, March 24). *OpenAI: ChatGPT payment data leak caused by open-source bug*. BleepingComputer. https://www.bleepingcomputer.com/news/security/openai-chatgpt-payment-data-leak-caused-by-open-source-bug/

Abril, D. (2023, March 14). Google is adding AI to its work apps. Here's what that means. *Washington Post*. https://www.washingtonpost.com/technology/2023/03/14/google-workspace-ai/

Agrawal, A., Gans, J., & Goldfarb, A. (2018, May 30). *Google's AI Assistant Is a Reminder that Privacy and Security Are Not the Same*. Harvard Business Review. https://hbr.org/2018/05/googles-ai-assistant-is-a-reminder-that-privacy-and-security-are-not-the-same

Auxier, B., Rainie, L., Anderson, M., Perrin, A., Kumar, M., & Turner, E. (2019, November 15). *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information*. Pew Research Center. https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/

Bolton, T., Dargahi, T., Belguith, S., Al-Rakhami, M. S., & Sodhro, A. H. (2021). On the Security and Privacy Challenges of Virtual Assistants. *Sensors*, *21*(7), 2312. https://doi.org/10.3390/s21072312

Chang, A. (2018, March 23). *The Facebook and Cambridge Analytica scandal, explained with a simple diagram*. Vox. https://www.vox.com/policy-and-politics/2018/3/23/17151916/facebook-cambridge-analytica-trump-diagram

Cho, E., Sundar, S. S., Abdullah, S., & Motalebi, N. (2020). Will Deleting History Make Alexa

More Trustworthy? *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3313831.3376551

Clauser, G. (2019, August 8). *Amazon's Alexa Never Stops Listening to You. Should You Worry?* Wirecutter: Reviews for the Real World; The New York Times. https://www.nytimes.com/wirecutter/blog/amazons-alexa-never-stops-listening-to-you/

D'Agostino, F., Gaus, G., & Thrasher, J. (2017). *Contemporary Approaches to the Social Contract (Stanford Encyclopedia of Philosophy)*. Stanford.edu. https://plato.stanford.edu/entries/contractarianism-contemporary/

February 11, C. N., 2020, & Am, 9:05. (n.d.). *Data from Equifax credit hack could "end up on the black market," expert warns*. Www.cbsnews.com. https://www.cbsnews.com/news/china-denies-responsibility-in-equifax-breach-after-doj-charges-four-military-members/

Fowler, G. (2019a, May 6). Perspective | Alexa has been eavesdropping on you this whole time. *The Washington Post*. https://www.washingtonpost.com/technology/2019/05/06/alexa-has-been-eavesdropping-you-this-whole-time/

Fowler, G. (2019b, May 28). It's the middle of the night. Do you know who your iPhone is talking to? *The Washington Post*. https://www.washingtonpost.com/technology/2019/05/28/its-middle-night-do-you-know-who-your-iphone-is-talking/

Gravrock, E. von . (2022, March 31). *Artificial intelligence design must prioritize data privacy*. World Economic Forum. https://www.weforum.org/agenda/2022/03/designing-artificial-intelligence-for-privacy/

Heikkilä, M. (2022, May 24). *The walls are closing in on Clearview AI*. MIT Technology

Review. https://www.technologyreview.com/2022/05/24/1052653/clearview-ai-data-privacy-uk/

Hern, A. (2019, August 2). *Apple contractors "regularly hear confidential details" on Siri recordings*. The Guardian; The Guardian. https://www.theguardian.com/technology/2019/jul/26/apple-contractors-regularly-hear-confidential-details-on-siri-recordings

Hill, K. (2020, January 18). The Secretive Company That Might End Privacy as We Know It. *The New York Times*. https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html

Johnson-Gomez, A. (2023, February 6). *A "Living" AI: How ChatGPT Raises Novel Data Privacy Issues | LawSci Forum*. LawSci Forum; University of Minnesota Law School. https://mjlst.lib.umn.edu/2023/02/06/a-living-ai-how-chatgpt-raises-novel-data-privacy-issues/

Kerry, C. F. (2020, February 10). *Protecting privacy in an AI-driven world*. Brookings. https://www.brookings.edu/research/protecting-privacy-in-an-ai-driven-world/

Kieren, M. (2018, December 20). *2018 ain't done yet... Amazon sent Alexa recordings of man and girlfriend to stranger*. Www.theregister.com. https://www.theregister.com/2018/12/20/amazon_alexa_recordings_stranger/

Kröger, J. L., Lutz, O. H.-M., & Raschke, P. (2020). Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference. *Privacy and Identity Management. Data for Better Living: AI and Privacy, 576*, 242–258. https://doi.org/10.1007/978-3-030-42504-3_16

Lapowsky, I. (2019, March 17). *How Cambridge Analytica Sparked the Great Privacy*

Awakening. WIRED; WIRED. https://www.wired.com/story/cambridge-analytica-facebook-privacy-awakening/

Lerman, R. (2021, September 2). Lawsuits say Siri and Google are listening, even when they're not supposed to. *Washington Post*. https://www.washingtonpost.com/technology/2021/09/02/apple-siri-lawsuit-privacy/

MacKinnon, E., & King, J. (2022, January 11). *Regulating AI Through Data Privacy*. Stanford HAI. https://hai.stanford.edu/news/regulating-ai-through-data-privacy

Medairy, B. (n.d.). *4 Ways to Preserve Privacy in Artificial Intelligence*. Www.boozallen.com. https://www.boozallen.com/s/solution/four-ways-to-preserve-privacy-in-ai.html

News, A. B. C. (2017, January 15). *Could Your Amazon Alexa Be Used Against You?* ABC News. https://abcnews.go.com/Business/amazon-alexa/story?id=44759936

Rouméas, É. (2023). The right to a fair exit. *Politics, Philosophy & Economics*, 1470594X2311569. https://doi.org/10.1177/1470594x231156939

Schonfeld, E. (2010, February 4). *Siri's IPhone App Puts A Personal Assistant In Your Pocket*. TechCrunch. https://techcrunch.com/2010/02/04/siri-iphone-personal-assistant/

Tiku, N., Vynck, G. D., & Oremus, W. (2023, January 27). Big Tech was moving cautiously on AI. Then came ChatGPT. *Washington Post*. https://www.washingtonpost.com/technology/2023/01/27/chatgpt-google-meta/

Valinsky, J. (2019). *Amazon reportedly employs thousands of people to listen to your Alexa conversations*. CNN. https://www.cnn.com/2019/04/11/tech/amazon-alexa-listening/index.html

Waddell, K. (2016, May 24). *Why Digital Assistants Are a Privacy Nightmare*. The Atlantic. https://www.theatlantic.com/technology/archive/2016/05/the-privacy-problem-with-

digital-assistants/483950/