Chapter 8
Privacy Issues of AI

This chapter sheds light on how private data is systematically collected, stored and analysed with the help of artificial intelligence. We discuss various forms of persistent surveillance at home and in public spaces. While massive data collection raises considerable ethical concerns, it is also the basis for better performance for AI systems.

8.1 What Is Privacy?

In its most basic form, privacy is the right to not be observed. People often act in a manner so as to increase their own privacy. For instance, shutting doors, wearing sun glasses, or clothing that is less revealing are simply subtle ways that humans actively moderate their own privacy.

Privacy is valuable for a number of important reasons: It allows people to make their own, non-coerced decisions, to better calculate their behaviour and be strategic in their social interactions, and also to take decisions and actions that do not conform to certain social norms.

Individual privacy has long had a contentious relationship with government. Just as individuals have typically fought for the right not to be observed, governments have often fought for the right to observe their own citizens.

Governments have long argued that "the need for privacy" argument may be used to cloak the planning and perpetration of crimes. Governments often argue that privacy must be curtailed in order to allow for proper law enforcement. Governments have argued that in some contexts, such as prisons, any expectation of privacy is not

C. Bartneck et al., *An Introduction to Ethics in Robotics and AI*,
SpringerBriefs in Ethics, https://doi.org/10.1007/978-3-030-51110-4_8

valid, but even in many public situations, most expectations of privacy need to be curtailed.

Governmental control over privacy has often been asserted to be a crime prevention method, but it can also serve the purpose of a general method of control. Limiting privacy provides governments with information about which individuals within a society need to be most controlled.

In the United States, questions about the right to privacy are fluid and subject to constant change as court decisions have, at times, awarded individuals more or less privacy. The United States Supreme Court has often considered privacy as a function of context, where a person's physical environment determines what privacy they should expect. But other considerations, such as the person's age or mental capacity may also be applied.

Companies and businesses are a third group that has traditionally been neutral in this debate. Companies try to perform a balancing act by weighing the needs and demands of customers against the needs and demands of government. Businesses often attempt to protect their own privacy and the privacy of their customers. Apple, for example, continues to deny government agencies systematic access to iPhones that have been locked, even if the data in the phone might help in the prevention or prosecution of crime (Holpuch 2016).

Companies may believe that if customers do not feel that their transactions with a company are private then they will no longer be customers. But companies must also rely on governments for a variety of policy, legal, and contractual obligations. Governments may use these powers as a method of persuasion to influence companies to hand over private data.

For centuries, this dynamic has existed, with groups not gaining much leverage over the other. Historically even the most authoritarian regimes had insufficient technological prowess to actually capture and keep track of private actions of their citizens. In the past 50 years this has changed. The development of computers and methods of mass data storage has allowed governments and business to collect, store, and process voluminous amounts of data, not only about its citizens but about all global citizens. Even more recently, the development of sensors capable of capturing data has broadened the horizon of where, when, how, and what data can be collected.

Finally, AI companies, long neutral in the debate over privacy, have begun to find value in the data that they have. They can now use AI methods to locate, model, and test potential products on populations of people. Using the psychological techniques discussed in Chap. 7, they can use powerful methods that exploit human tendencies for the purposes of improved advertising, marketing, and sales. The result is that the right to privacy is eroding. So is there any need for these data? Or is data collection simply done for its own sake?

8.2 Why AI Needs Data

Current robots and AI systems collect enormous quantities of data about their users. These data collections require considerable effort, and it might not be obvious at first sight what the benefit of this collection epidemic is.

Let’s consider the case of Facebook. Mark Zuckerberg was summoned to appear before a Senate Judiciary and Commerce Committee hearing on April 10, 2018. During the hearing Senator Orrin Hatch asked: “How do you sustain a business model in which users don’t pay for your service?”, to which Zuckerberg simply responded: “Senator, we run ads.”.

John Wanamaker famously pointed out that “Half the money I spend on advertising is wasted; the trouble is I don’t know which half.” What makes Facebook such an attractive platform for placing advertisement is that it knows a lot about its users and allows marketing experts to target specific groups. This targeting dramatically improves the efficiency and effectiveness of the ads. Less than the famous 50% of the advertisement gets wasted this way. The German online retailer Otto, for example, was able to reduce the cost per click by 38% through highly targeted advertisements (Kleinz 2018). Facebook users that previously ‘liked’ a picture of a friend in Hawaii were presented with advertisements for cheap flights to that destination.

While the advertising industry might be one of the greatest beneficiaries of private data, its use stretches far beyond this. A companion robot is likely to be perceived as more relatable and friendly if it knows and adapts to its specific user. A smart home might also learn when its inhabitants are present and adapt the heating to maximise comfort and energy efficiency. In summary, the more and better information about users is available, the better AI systems can learn from the data and adapt their behaviour.

8.3 Private Data Collection and Its Dangers

Gathering personal data has become dramatically easier with the arrival of certain key technologies, such as smartphones, surveillance cameras and of course, the Internet. These days it is in principle possible to track every step users take and every restaurant they visit. People take photos of the food they eat and post them online. Within the framework of the Self Optimisation movement people feverishly collect personal data in an attempt to change their lives for the better. Much of this data is now being uploaded to cloud computers, which has significantly increased the possibility for tracking private information.

Moreover, users of social networks voluntarily upload very private data and seem to deliberately ignore that by uploading data they often transfer the copyright of this data to the platform provider. Facebook and others *own* the data and use it and even sell it to others.

One of the factors of Google’s success in gathering personal data is that people cannot hide their interest when searching for information. While many would try to conceal delicate private issues they cannot search for information about this topic without entering the terms into the search box. Stephens-Davidowitz and Pinker (2017) analysed such personal search query patterns and identified a considerable group of Indian husbands that desire to be breast-fed. How Google will respond to

Fig. 8.1 Amazon Echo Plus uses Alexa (*Source* Amazon)

these results remains to be seen, but it does show that even our most intimate wishes and issues are being collected online.

Autonomous vehicles typically also report back telemetry data about the cars' performance, which in turn allows these companies to compile reports on a car's safety. Tesla, for example, compiles a quarterly report on the kilometres driven by its vehicles and whether the auto pilot was engaged.

8.3.1 Persistence Surveillance

Persistent surveillance is the constant observation of a person, place or thing. Within the military and policing fields, persistent surveillance is a commonplace technique for gathering information about an enemy or suspect. Yet, with the development of so-called digital assistants such as Amazon's Alexa (see Fig. 8.1) and Google Home, similar elements have become a product feature.

These systems stream audio data from the home to the parent company where the data is stored, collected, and analysed. Not only is this data quickly examined for requests from the product, but it might theoretically also be used to observe users and their environment in ways that are unknown to them. Sensitive to these concerns, Amazon has put a number of features in place to limit the data collecting ability of its devices. According to Amazon, features such as using the word “Alexa” to wake the device up prevents the device from being used as a means for persistent surveillance (Hildrenbrand 2018). Yet, as shown through the Wikileaks release of NSA documents, backdoors and vulnerabilities can be exploited in these types of technologies which could theoretically convert them into a means of persistent surveillance.

Another prominent example of private data collection is the “Hello Barbie” doll built by Mattel. It uses cloud-based speech recognition, conversation management and speech synthesis. While this is not unlike Alexa or Google Home, “Hello Barbie” was targeted at young girls. They were encouraged to talk with Barbie about their lives, and since they might not have been aware of the technical and social implications, they would likely have shared unfiltered thoughts with Mattel. Of course, Mattel vowed to respect the privacy of its customers, but it remains a question of trust whether parents believed them (Fig. 8.2).

Consumers, on the other hand, seem somewhat ambivalent towards privacy concerns about these devices. In a study conducted by Shields (Shields 2018), for instance, consumers report that they like the home security features of a digital assistant, yet also fear being spied on. Privacy-related events can cause citizens to catch on to privacy-related issues. A study by Anton et al. (2010) found that internet users had become more concerned about privacy even as their usage increased from 2002 to 2008.

In contrast to digital assistants which are willingly placed in the home by a person, persistent surveillance can also be conducted from a distance. Drone surveillance developed for the battlefield allows for continuous observation of individuals, i.e., the collection of information about individuals including the creation of social networks, and the creation of behavioural models that capture patterns of behaviour. In warfare, these social networks are used to distinguish non-combatants from enemy soldiers and to predict upcoming attacks. Recently, these practices have been adapted from the battlefield and applied to domestic locales. In 2005, Baltimore created a ground level surveillance system called CitiWatch which contained more than 700 cameras placed around the city.

In October 2014, CitiWatch expanded to create a database which privately owned surveillance cameras could contribute to on a voluntary basis. Later on, the Baltimore program began to include aerial surveillance from camera-laden air planes. On the one hand, Baltimore law enforcement officials claim that these programs work to reduce crime. On the other hand, privacy advocates claim that these programs greatly restrict privacy and inhibit personal freedom. Informal polls conducted by the Baltimore Business Journal and the Baltimore Sun found that 82% of people felt “comfortable” with the surveillance program “as long as it’s keeping people safe.” Moreover, an online Baltimore Sun poll found that 79% of people believed that the police department’s level of secrecy around the program was acceptable (National Police Foundation 2017).

Similarly, airline authorities have considered continuously streaming the audio from airplane cockpits. Doing so would reduce the need for retrieving black box recorders. But pilots have hitherto resisted streaming audio from the cockpit on the

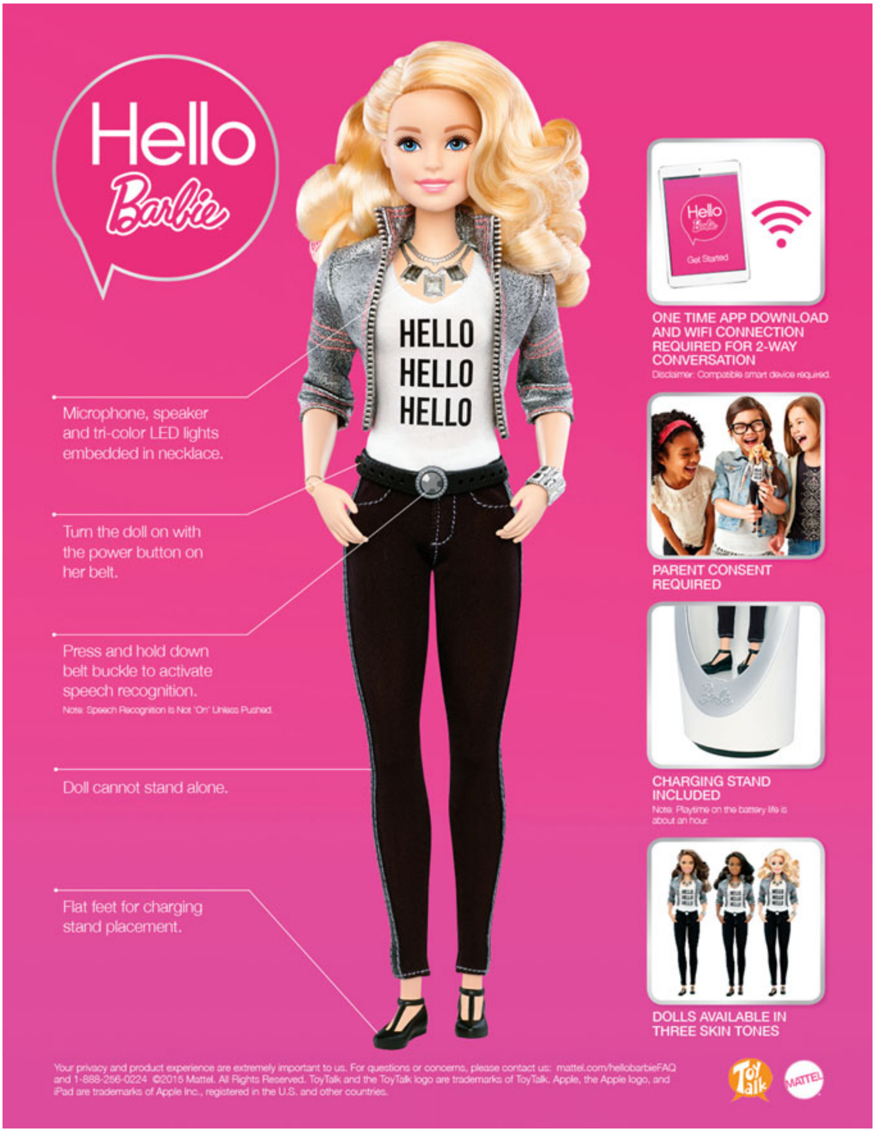

Microphone, speaker and tri-color LED lights embedded in necklace.

Turn the doll on with the power button on her belt.

Press and hold down belt buckle to activate speech recognition.
Note: Speech Recognition Is Not 'On' Unless Pushed.

Doll cannot stand alone.

Flat feet for charging stand placement.

ONE TIME APP DOWNLOAD AND WIFI CONNECTION REQUIRED FOR 2-WAY CONVERSATION
Disclaimer: Compatible smart device required.

PARENT CONSENT REQUIRED

CHARGING STAND INCLUDED
Note: Playtime on the battery life is about an hour.

DOLLS AVAILABLE IN THREE SKIN TONES

Your privacy and product experience are extremely important to us. For questions or concerns, please contact us: mattel.com/hellobarbieFAQ and 1-888-256-0224 ©2015 Mattel. All Rights Reserved. ToyTalk and the ToyTalk logo are trademarks of ToyTalk. Apple, the Apple logo, and iPad are trademarks of Apple Inc., registered in the U.S. and other countries.

Fig. 8.2 Hello Barbie (*Source* Mattel)

basis that doing so is an invasion of privacy. There is still a lot of uncertainty about how to manage human expectations in environments that feel as if they offer privacy, when actually they might not.

Hence we see that even large parts of a population may be willing to sacrifice their privacy if they see some benefit. Many cities around the world now contain cameras and systems in place for persistent surveillance, yet at least a number of countries seem to be willing to accept this. The use of unmanned aerial vehicles (UAVs) for the purpose of surveillance can also raise privacy issues. These systems have been used by individuals to spy on their neighbours. Regulations related to drone use vary from none at all to a complete prohibition on drone imports and use. The European Union, for example, has strong privacy protections in place which extend to data collected using a drone.

8.3.2 Usage of Private Data for Non-intended Purposes

While the availability of users' private data enables AI systems to perform better, there are also considerable risks associated with this data collection. One of the main issues is the usage of data for non-intended purposes. Users are often unaware how their data will be processed, used and even sold.

The success of programmatic advertisement demonstrates how personal data can be used to make people buy products. This specific form of marketing has in general been accepted by society, and people have developed protection mechanisms. We know that advertisements can only be trusted to some degree. But we have not yet developed such a critical awareness towards AI. We might not even be aware of manipulation taking place.

But there are serious forms of manipulation that use the exact same mechanisms that marketing experts applied so successfully: AI systems and social robots can be used to manipulate opinions on anything, including political views. Chapter 7 talks in more detail about the influence of AI on the public discourse. The Cambridge Analytica scandal in 2018 demonstrated how private data gathered through Facebook can be used in attempts to manipulate elections.

From a privacy perspective, the ability of modern AI to impersonate people has become a true danger. In 2016, Adobe demonstrated their VoCo system that can imitate the voice of any speaker after listening to approximately 20 min of conversation. But AI did not stop at speech. The manipulation of video, in particular the exchange of faces, has become a major issue that takes the problems of revenge porn to a new level. Spiteful ex-partners are able to mount the face of their previous love onto actors in porn movies and share them online. These so called "deep fakes" use neural networks to manipulate videos (see Fig. 8.3).

Another problem with the large scale collection of private data is that many people are unaware of the contracts they enter when signing up for various online services. The Terms and Conditions of social networks, for example, are not an easy read. The consequences, however, are not to be underestimated. As mentioned above, Facebook owns all the photos, messages and videos uploaded, and so does Google. And they do not shy away from selling this data to others. The General Data Protection Regulation (GDPR) in effect in the European Union since 2018 forced companies to seek consent from users before sharing their personal data with others. It should however also be noted that critics of the GDPR say that it leads to counterproductive results, by making small organisations and individuals shut down their websites for

Fig. 8.3 Former US president Barack Obama was used to showcase the power of deep fakes (*Source* BuzzFeedVideo)

fear of being heavily fined, while large companies can easily deal with the additional workload of complying with the new regulation.

Vulnerable Populations

Privacy violations and limitations can have different effects on vulnerable populations. Data generated by those receiving medical attention generates additional privacy concerns. This data, for instance, can be used not only to understand and treat an individual's affliction, but can also be used to infer things about the person's family genetics, physical and mental limitations, and perhaps even predict their death (White et al. 2012). Additional laws and regulations exist for health related data.

Artificially intelligent technologies foster dilemmas by generating data that would otherwise be private or using data in ways that were previously not possible. For example, gait (movement) information can be passively captured by camera observations of the person. This information could theoretically, in some cases, be used to predict an older adult's mortality. In a worst-case scenario, a private individual or business could passively observe large numbers of older adults, predicting their impending mortality, and then attempt to sell them funeral packages based on their data.

Children represent another vulnerable population. As mentioned above, data can be used to directly market products ranging from breakfast cereals to toys for children. Moreover, data generated when the child plays with an artificially intelligent toy can be transmitted to the company of origin and used for profiling, marketing, and advertising purposes. In general, parents tend to be more concerned about privacy when being observed with their children.

8.3.3 Auto Insurance Discrimination

Lack of privacy and the amount of data collected may lead to different rules being applied to different groups. Insurance companies, for example, value data in order to predict the cost of a policy and to assign premiums. Auto insurance companies use data to evaluate the driving behaviour of their drivers and to assign them to risk classes. AI use could lead to bias and discrimination here.

8.3.4 The Chinese Social Credit System

The Chinese government started work on a Social Credit System in 2014 that collects vast amounts of information about its citizens. While credit rating agencies have been operating for far longer, the Chinese government intends to extend the reach of its data collection far beyond what other organisations typically cover. The Chinese Social Credit System is already operational to the level of providing a financial credit score. In the future, it is intended to consider more private data, such as web browsing behaviour, and calculate how good a citizen is. This would have much further consequences than the denial of credit. Chinese citizens with a low score might be banned from flying and excluded from private schools, hotels and even careers.

Given the centralised authority of the Chinese government, its technical abilities and its plans for the future the Social Credit system could become the archetype for a tight control of the collection and use of private data. Its potential for use and abuse are still unclear. The opaque nature of the current system and the lack of a free press make it vulnerable to systematic bias.

8.4 Future Perspectives

From a technical perspective, it is relatively easy to apply the principles of programmatic advertisement to other application areas. Autonomous vehicles could, for example, make real-time bids for the lives of their passengers based on their social credit score in the face of an imminent crash. The passengers with the higher score would less likely to be harmed, while passengers with a lower social credit score would be exposed to greater risks of harm. The technical problems could be solved and only our societal discourse could prevent such a dystopian future.

It is also of concern that many have become ambivalent to how their private data is being shared, used and sold. Not long ago, people would have hesitated sharing a photo of themselves with the whole world. Today, social influencers flood the internet with revealing photos.

The future seems to be moving towards greater data collection. The value placed on privacy depends a great deal on the culturally and historical tendencies of the populace. Artificial intelligence expands the ways in which data is used, and may offer some predictive ability with respect to what is being collected. Future generations will need to decide the boundary that defines privacy and data use.

Discussion Questions:

- What information would you not like your digital assistant to have about you? Create a list.
- What parts of the population are most vulnerable to an invasion of their privacy by an AI? Explain your reasoning.
- Do you think a social credit system would have benefits? Discuss.

Further Reading:

- Michael J Quinn. *Ethics for the information age (7th edition)*. Addison-Wesley Publishing Company, 2017. ISBN 978-0134296548. URL http://www.worldcat.org/oclc/1014043739 (Chap. 5)
- Sare Baase and Timothy M. Henry. *A Gift of Fire Social, Legal, and Ethical Issues for Computing Technology (5th Edition)*. Prentice Hall PTR, 2017. ISBN 9780134615271. URL http://www.worldcat.org/oclc/1050275090 (Chap. 2).