Data ethics and privacy: should, can, will?

When using technology such as artificial intelligence (AI), ethical considerations play a major role in our society. There is a reason for this: as we increasingly face public scandals related to the misuse of personal data, the call for responsible policies concerning ethics and privacy is growing. The trust that customers, employees and citizens have in both public and private organizations is at stake. The critical question for organizations is: how do we get the most out of what data and technology have to offer while simultaneously addressing ethical and privacy concerns?

This article takes a closer look at data ethics and the intersections with privacy, discusses the legal developments and provides practical tips on how to get started setting up and strengthening data ethics in organizations.

Marijn Bakker LL.M., CIPM, CIPP/E,
is a data privacy manager at KPMG Cyber & Privacy.

Sterre Stolk LL.M.
is a former privacy consultant at KPMG Netherlands.

INTRODUCTION

May 25, 2023, marks the fifth anniversary of the European privacy law, the General Data Protection Regulation (GDPR). For many organizations, privacy protection is now an integral part of their business operations. However, there is still more that can be done.

Even with the introduction of the GDPR, Dutch people's confidence that companies and government organizations handle their privacy well remains low. A 2021 privacy survey of a sample of the Dutch population ([KPMG21]) showed that a quarter of Dutch people harbor considerable distrust of the government, and their trust in technology companies is even lower. This manifests itself in growing concerns about their own privacy.

Whereas trust in government agencies and companies is declining, interest in privacy is increasing. An overwhelming majority of the Dutch (86 percent) think it is good that there is a lot of focus on privacy ([KPMG21]). This is substantially more than at the beginning of 2018, when the KPMG survey "A little privacy please" showed that 69 percent considered privacy an important issue ([KPMG18]; see also [ICTM18]). In addition, this interest is confirmed by the fact that the Netherlands is one of the leaders within the European Union in terms of reporting data breaches ([DLAP21]). One explanation for this increasing attention to privacy is the continuing developments in the digital transformation that society is undergoing. As a result, this question is now at the forefront of privacy and data ethics debates: how responsible or ethical are all the technical developments that succeed one another in a relatively short period of time?

The Dutch Data Protection Authority, the Autoriteit Persoonsgegevens (AP), concluded in its annual report for the year 2021 that society has reached the point where digitization can no longer take place without ethical value considerations ([AP22b]). In their private lives, as consumers and citizens, people are constantly confronted with new technologies, although they may not always realize it. Technologies such as algorithms have taken a structural place in our daily lives. Whether it is making a purchase in an online store and paying afterwards, taking out a loan from a bank, or applying for a grant from the government, there is a good chance that the application will be assessed using technology.

New technologies bring tremendous opportunities for developing new products and services, ensuring a better customer experience, and improving efficiency in the workplace. However, to ultimately continue the successful integration of new technologies, organizations must use them responsibly. Data ethics and privacy play an essential role in this regard. The GDPR provides guid-

Figure 1. The trust Dutch citizens have in their government is not very high, but the trust they have in large technology companies appears to be even lower ([KPMG21]).

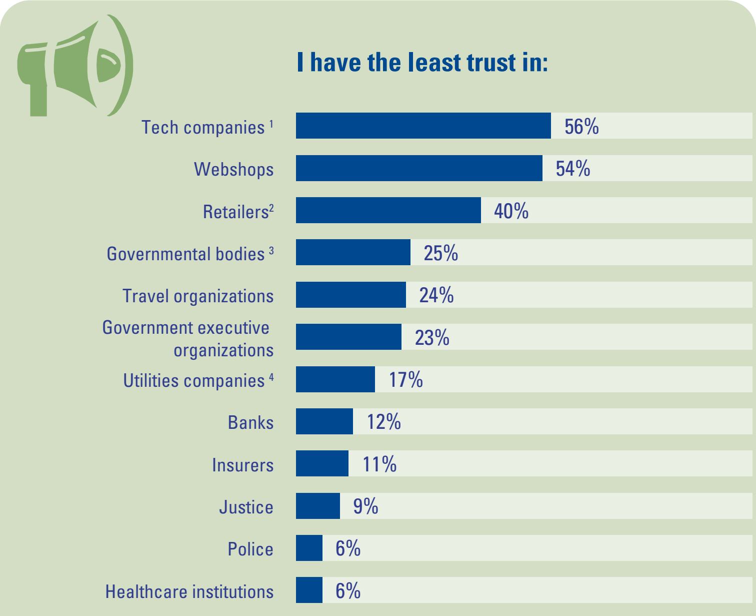

[1] For example, Apple, Google Facebook
[2] For example, supermarket, clothing shops, etc.
[3] For example, ministries and local governments
[4] For example, energy companies and network operators

ance on the ethical use of personal data, but data ethics is broader than just the collection, use, retention and deletion of personal data.

WHAT IS DATA ETHICS?

Ethics is about right and wrong. About what society, citizens or consumers think is fair, just and acceptable ([Meel]). Viewed from a privacy perspective, data ethics is not so much about whether an organization may process personal data and whether the processing meets the requirements of the GDPR, but rather it is about a question that is more fundamental. Even if organizations *could* or *want* to do something (e.g., from both a legal or technological perspective), organizations must continually ask themselves whether they *should* from an ethical perspective. In other words: it is allowed, but is it the right thing to do?

Data ethics requires a different way of thinking within an organization that focuses on the impact a data operation has on people and society. Data ethics revolves around this question: from an ethical perspective, is what we want to do or have the capabilities for, the right thing to do? A privacy professional may see common ground with conducting a Data Protection Impact Assessment (DPIA), which identifies the privacy risks of a personal data processing activity to data subjects. However, data ethics is much broader than privacy. Data ethics is about non-discrimination, avoiding bias and acting transparently and responsibly toward people affected by the use of technology. The example in the box on the right illustrates this ([Ober19]).

A hospital in the United States used a commercial algorithm to determine which group of patients might need more care than average. Algorithms are widely used in the United States by hospitals, government agencies, and other healthcare providers. It is estimated that about 200 million people are assessed annually using such algorithms.

In this case study, an algorithm ranked patients based on the extent to which they used medical care. Patients in the 97th percentile and above were marked as "high risk" by the algorithm and were automatically enrolled in a health program. Wanting to improve the health of high-risk patients is a noble goal, but in retrospect it was found that a bias based on race was present in the algorithm. According to the algorithm, patients of color are healthier than white patients, but this turned out to be the wrong conclusion.

The reason for this bias could be traced to the input data. People of color are less likely to use healthcare services than white people and spend an average of $1,800 less per year on health care than white people. The algorithm inferred that people of color must be healthier, since they use fewer healthcare services. However, this assumption was incorrect. The dataset on which the algorithm was based, consisted of 44,000 white patients and only 6,000 patients of color. Because of this limited input data, the algorithm made incorrect assumptions that had a negative impact on healthcare access for a certain group of people.

LEGAL DEVELOPMENTS IN DATA ETHICS: THE AI ACT

When it comes to data ethics in the public debate, it is regularly about the reliability and the increasing use of data and algorithms in private and public organizations, and also about how to use algorithms in a controlled, ethical way. Worldwide, the European Commission has taken the lead in regulating the use of artificial intelligence. This has resulted in a bill called the Artificial Intelligence Act (AI Act).

The AI Act aims to establish a framework for responsible AI use. According to the Act, AI systems must be legally, ethically and technically robust and must respect democratic values, human rights and the rule of law. The fact that the AI Act is focusing on regulating the use of AI from a technical and legal perspective, is not surprising. What is unique to this Act, is the strong emphasis on data ethics. The aim is to reach a final agreement on the AI Act this year (2023), but there is no concrete deadline. When a final agreement is made, there will be a grace period of around two years to allow affected parties to comply with the regulations.

The AI Act introduces new obligations for companies and governments, as well as a supervisory authority and a penalty system. These are detailed in the sections below. It is important to emphasize that no final agreement has been reached on the exact content of the AI Act. In other words, legal developments (and proposed amendments to the AI Act[1]) are rapidly following one another. For example, adjustments to the AI Act are currently being con-

[1] The final content of the AI Act is currently still being negotiated in Europe. This means that this article provides an insight into the developments concerning the AI Act but cannot provide certainty on the final content of the AI Act.

sidered to deal with developments around ChatGPT and similar generative AI models. In other words, the legal and technical developments that may have an impact on the AI Act, are worth keeping an eye on.

Conformity assessment for high-risk AI systems

The AI Act introduces conducting a so-called conformity assessment by an outside body. In other words, if an AI system could pose a high risk to the health, safety or fundamental rights of people, its providers must have an assessment conducted by an independent third party to identify and mitigate those risks. These assessments help ensure compliance with the AI Act. For AI systems with limited or minimal risk, less onerous requirements apply. In that case, a self-assessment or transparency requirement is sufficient.

The legislative proposal for the AI Act currently states that the European Commission is the body that determines what constitutes a high-risk system and when a mandatory conformity assessment must be conducted. AI systems that will potentially qualify as high risk include systems for migration and asylum, critical infrastucture, law enforcement, and product safety. In addition, it is currently being examined whether generative AI models such as ChatGPT should also be regarded as high risk.

Based on the proposed AI Act, there is also the possibility that an AI system classifies as a high-risk system, but a conformity assessment is not required. In such cases, a self-assessment is sufficient. Currently, the AI Act states that the European Commission will determine for which (high-risk) AI systems a self-assessment should be performed.

High-risk AI systems must meet strict requirements under the AI Act before they can be marketed. Measures to be implemented under the proposed AI Act include: establishing a risk management process that specifically oversees the AI application, setting high data quality requirements to prevent discrimination, maintaining logs, establishing documentation around accountability, ensuring transparency, establishing a system in which people oversee the AI applications, and ensuring security and accuracy standards.

AI database for high-risk systems

Another new aspect of the AI Act relates to the creation of an AI database in which high-risk AI systems are to be registered. The AI Act currently states that the database will be managed by the European Commission and aims to increase transparency and facilitate oversight by regulators.

Introduction of national AI supervisor

The proposed AI Act currently contains an obligation for each member state to form or designate an authority to supervise compliance with the AI Act. This national supervisory authority will participate in the European AI Board (EAIB), which will be chaired by the European Commission and will also include the European Data Protection Supervisor (EDPS). Recently, the Dutch Data Protection Authority, the AP, was appointed as algorithm supervisor in the Netherlands. With this appointment, the Netherlands is already fulfilling its future obligation under the AI Act.

Fines for failure to comply with AI Act

Like the GDPR, the AI Act will include a penalty system. The biggest fine that can be imposed under the Act is a fine of up to 30 million euros or 6 percent of annual global turnover, whichever is higher. This is 2 percent higher than the highest fine category under the GDPR. Aside from the moral obligation for companies to take data ethics and privacy seriously, there will be financial incentives to set up AI systems in accordance with the upcoming AI Act.

HOW TO PUT DATA ETHICS INTO PRACTICE

It is clear that the AI Act will make its appearance in the future. Nevertheless, it is important to realize that legislation is only the basis and that acting ethically requires more than complying with new legislation. What can organizations do today to promote the ethical handling of data and raise awareness in their organization?

Contact the Privacy Officer

The concerns that exist about AI systems are in many cases about the use of personal data. Despite the fact that privacy and data ethics are two different topics, they often overlap. This means that if an organization has appointed a Privacy Officer, in all likelihood they are already working on the topic of data ethics in the light of the use of personal data.

The GDPR has an obligation to conduct a DPIA on personal data processing activities that may result in a high privacy risk. In many cases, this obligation will also apply to AI systems that process personal data. Even though the AI Act focuses on the quality of AI systems while the GDPR focuses on the use of personal data, the two laws converge when personal data is used in AI systems. Therefore, Privacy Officers can be a good starting point to prepare the organization for the upcoming AI Act. They can help identify which systems in the organization use AI and whether these systems may pose a high risk.

Establish an ethical framework

The first step to securing data ethics in an organization is to establish what ethical handling of data specifically means for the organization. This can be done by formulating values and principles around the topic of data ethics, for example an ethical framework or compass. It is important that the ethical principles align well with the culture and core values of the organization and are recognizable to employees from all levels of the organization ([Meel]).

Organize independent oversight

Data ethics is an abstract topic, but it needs a very concrete interpretation. Most organizations are not (yet) equipped to deal with the ethical dilemmas that arise when new technologies, such as algorithms, are deployed. Furthermore, there is often no monitoring of the integration of ethical principles into business operations. A powerful tool in both establishing ethical principles and closing the gap between principles and practice, is the establishment of effective and independent oversight. This can be done by an independent

According to FRAIA, the decision-making process regarding algorithms can be divided into three main stages:

- Stage 1: preparation. This stage is about deciding why an algorithm will be used and what its effects will be.
- Stage 2: input and throughput. This stage is about the development of an algorithmic system. In this stage, it is decided what the algorithm must look like, and which data is being used to feed the algorithm. Within this stage, the FRAIA further distinguishes between:
  - Stage 2a: data, or input. This involves asking questions that pivot on the use of specific data and data sources.
  - Stage 2b: algorithm, or throughput. This involves questions regarding the algorithm, and its operation and transparency.
- Stage 3: output, implementation and supervision. This stage is about how to use the algorithm, i.e., about the question which output the algorithm generates, how that may play a role in policy or decision-making, and how that can be supervised.

Source: [MBZK21]

> Ethics is about right and wrong. About what society, citizens or consumers think is fair, just and acceptable

committee, internal audit teams, or an independent third party ([Meel]).

Conduct a Fundamental Rights and Algorithm Assessment

When an organization works with algorithms, it is wise not to wait for the introduction of the AI Act and to already start identifying risks when using algorithms. This can be done by conducting a Fundamental Rights and Algorithm Impact Assessment (FRAIA). FRAIA is the English translation of the Dutch "Impact Assessment Mensenrechten en Algoritmes" (IAMA). The FRAIA was developed by the Utrecht Data School and helps to make careful decisions about the use of algorithms. The FRAIA is mandatory for government agencies and can also help other organizations gain a better understanding of the considerations and risks involved in the decision-making process concerning the deployment of algorithms. It is also a way to "practice" the impending assessments that the AI Act will most likely introduce.

CONCLUSION

There is currently no clear body of standards, laws, or case law in the field of data ethics. While the AI Act aims to fill this gap, ethical handling of data requires more than following the letter of the law. Take the example of the GDPR, Europe's data privacy law. The GDPR gives us more control over our personal data, but the ethical principle of privacy is a much broader and abstract issue than simply protecting data. Therefore, an organization that sees its customers' privacy as its responsibility, will have to think beyond just avoiding a GDPR, and soon, an AI Act fine ([Meel]).

References

[AP22a] Autoriteit Persoonsgegevens (2022, March 15). *AP Inzet Artificial Intelligence Act*. Retrieved from: https://autoriteitpersoonsgegevens.nl/sites/default/files/atoms/files/ap_inzet_ai_act.pdf

[AP22b] Autoriteit Persoonsgegevens (2021). *Jaarverslag 2021*. Retrieved from: https://www.autoriteitpersoonsgegevens.nl/sites/default/files/atoms/files/jaarverslag_ap_2021.pdf

[DLAP21] DLA Piper (2021, January 19). *Nederland tweede van Europa in aantal gemelde datalekken sinds invoering AVG*. Retrieved from: https://www.dlapiper.com/en-nl/news/2021/01/nederland-tweede-van-europa-in-aantal-gemelde-datalekken-sinds-invoering-avg

[EC21] European Commission (2021). *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*. Retrieved from: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206

[ICTM18] ICT-Magazine (2018, October 16). *Onderzoek KPMG: Nederlander nauwelijks bekend met nieuwe privacyrechten*. Retrieved from: https://www.ictmagazine.nl/onderzoek-kpmg-nederlander-nauwelijks-bekend-met-nieuwe-privacyrechten/

[Meel] Van Meel, M. & Remmits, Y. (n.d.). *Risico's van algoritmes en toenemende vraag naar ethiek: Deel 4 - De burger en klant centraal bij het gebruik van algoritmes* [KPMG Blog]. Retrieved from: https://home.kpmg/nl/nl/home/topics/artificial-intelligence/vertrouwen-in-algoritmes/risicos-van-algoritmes-en-toenemende-vraag-naar-ethiek.html

[KPMG18] KPMG (2018). *Een beetje privacy graag*. [Report can be requested at KPMG.]

[KPMG21] KPMG (2021, October). *Meer zorgen over privacy: Het resultaat van ons privacy onderzoek onder consumenten*. Retrieved from: https://assets.kpmg/content/dam/kpmg/nl/pdf/2021/services/meer-zorgen-over-privacy-whitepaper.pdf

[MBZK21] Ministerie van Binnenlandse Zaken en Koninkrijksrelaties (2021, July). *Impact Assessment Mensenrechten en Algoritmes*. Retrieved from: https://open.overheid.nl/documenten/ronl-c3d7fe94-9c62-493f-b858-f56b5e246a94/pdf

[Ober19] Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019, October 25). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, *366*(6464), 447-453, Retrieved from https://www.science.org/doi/10.1126/science.aax2342

About the authors

Marijn Bakker is a manager at KPMG Cyber & Privacy. She has broad experience in GDPR implementation and optimization projects, privacy audits and assessments, and the implementation of privacy tooling such as OneTrust. Marijn holds a master's degree in Internet, Intellectual Property and IT Law from the Vrije Universiteit Amsterdam and a master's degree in Corporate Law from Leiden University.

Sterre Stolk is a former privacy consultant at KPMG Netherlands and now works as a senior privacy advisor in the public sector. At KPMG she worked on various privacy engagements mostly related to complex technological challenges. Sterre holds a master's degree in Internet, Intellectual Property and IT Law from the Vrije Universiteit Amsterdam.