Jeroen van Helden - Associate at De Clercq Advocaten Notariaat

Privacy-law aspects of Artificial Intelligence

We learn by making mistakes, is how the saying goes. This applies for a self-learning AI application as well. The logic of a self-learning AI application is not pre-programmed, but is given shape within an artificial neural network that is fed with examples, from which the system then learns. These examples can contain privacy-sensitive information. At a certain point, the application has learned enough; in other words, it is 'mature' enough to go 'live' and can actually start playing a role in the real world. The application may be used to operate a car, for instance, in which new data are in turn collected. Certain parties involved in the development of self-driving transport, such as Alphabet with subsidiary Waymo (Google), are suspected of ultimately being less interested in selling self-driving cars than in collecting the data that are generated about the driver and cashing in on those data. An AI application can also be used to take decisions that have a certain impact on a person's private life, for the selection of potential defrauders, for instance, or in screening job applicants. So it is high time to examine the privacy-law aspects of self-learning applications.

Framework

Any discussion of privacy will quickly arrive at a mention of the General Data Protection Regulation (GDPR). The GDPR took effect throughout the European Union on 25 May 2018 and provides rules for the (automated) processing of personal data. This type of privacy regulation arose in the 1970s when it became clear how easy it was to store and further process large volumes of data using computers. The GDPR is meanwhile quickly becoming the international standard for regulations on data processing.[1]

Because of the attention garnered by the GDPR, one might think that this European law is the only privacy-law framework which society needs to deal with in the development and application of artificial intelligence. That is not entirely correct. Privacy as a fundamental right has long been anchored in numerous international treaties, as well as in the Dutch Constitution. The fundamental right to the protection of our private life encompasses much more than the rules on processing personal data and can emphatically also be relevant for AI applications. This article first discusses AI in relation to the GDPR, before discussing AI in relation to the fundamental right to privacy.

Datasets

Whether an AI application is trained on the basis of supervised learning, unsupervised learning, reinforcement learning or some other learning method, data are always needed to train and develop the automated system, so that it can subsequently go to work independently. The GDPR interprets the term 'personal data' broadly, so that the GDPR will apply to many datasets.[2] This presents the following challenges:

- Based on the GDPR, personal data can only be collected for well-defined, explicitly described and justified purposes and the data may not be subsequently further processed in a manner inconsistent with those purposes. Data subjects must also be informed about that use and their permission may be required. The data system behind an AI application is often complex and seldom transparent. Aggregate and enriched data originating from multiple sources are used. In practice, the obligation of purpose limitation and duty to inform are therefore perceived to be difficult or even impossible to fulfil.[3]
- According to the GDPR, personal data must be

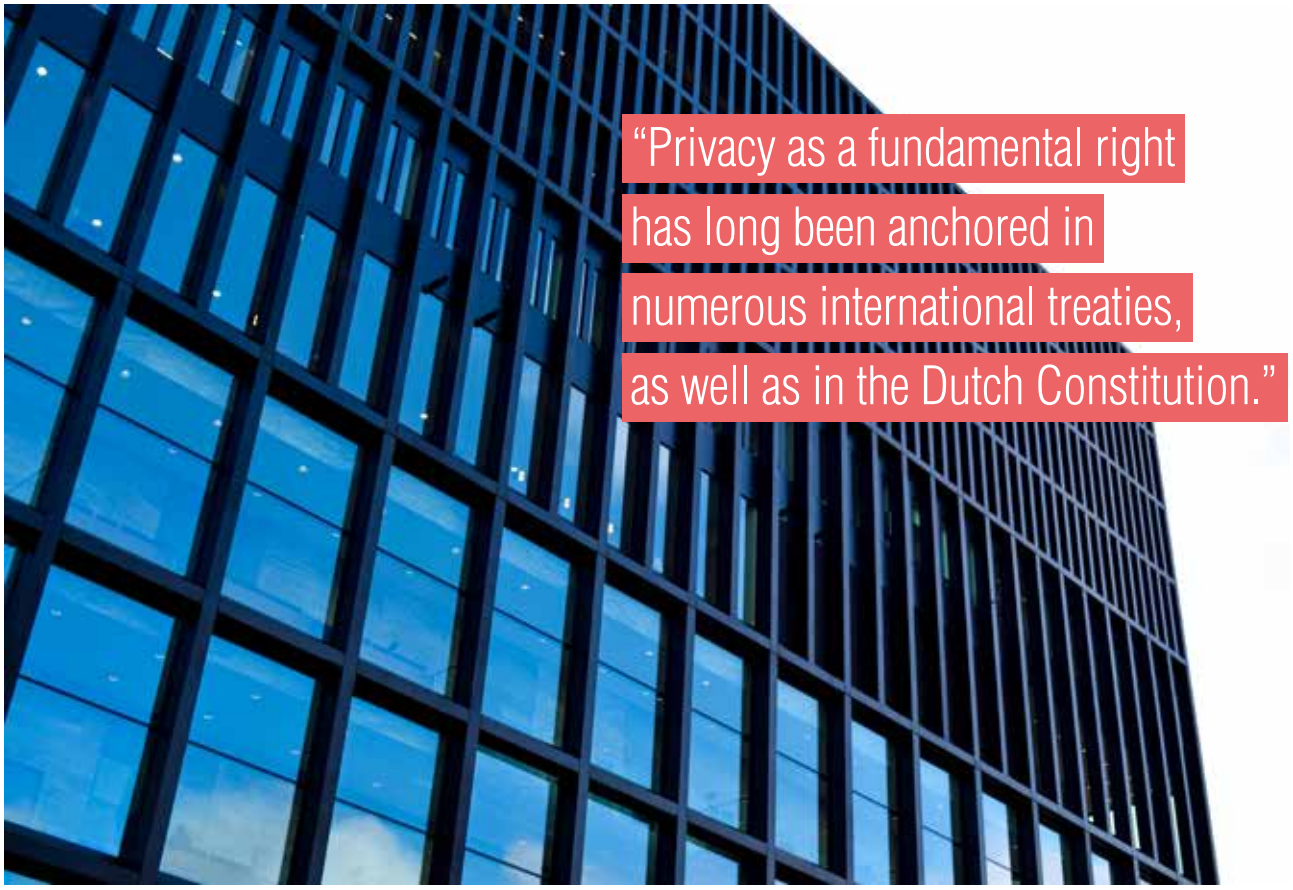

sufficient, fit for purpose and necessary for the purposes for which they are processed. The development of AI applications often benefits from large datasets. The trick is maximal data processing rather than minimal data processing. Although some techniques, like generating synthetic data or using Federated Machine Learning, can help limit the processing of personal data, there is no simple solution that eliminates the fundamental tension between the principles of the GDPR (few data) and AI applications (many data).[4]

When using AI applications, one must also be aware that AI applications are capable of discerning patterns in the dataset that humans would not be able to discover. As a result, ‘innocent’ data can become ‘sensitive’ data in the hands of an AI application. For instance, an AI application could be able to conclude a person’s political leaning or medical condition from what appear to be innocent data. It is also conceivable that an AI application could convert what appears to be anonymous information into data identifying an individual by name. Recent research also shows that self-learning AI applications do not quickly forget rare and sensitive training data and that they can unintentionally produce these data in live environments.[5]

Automated decision-making

Automated decisions are being taken all around us. For example, the question of which advertisement you are shown online or the amount of the supplement to which you are entitled according to the tax authorities’ website. The advantages are unmistakable: speed and consistency. It is suspected that most of the decisions in the sense of the General Administrative Law Act (Awb) are taken via an automated process.[6]

The GDPR contains a few specific rules on automated decision-making. First of all, an individual must be able to know that automated decision-making is being used. That is why the controller has the obligation to inform the data subject about this.[7] Part of this duty to inform is that the data subject must be informed about the ‘underlying logic’ of the application. This does not require that the algorithm’s functioning be explained in (technical) detail or that the algorithm be published. It must be made clear in a comprehensible manner, however, how

"One must also be aware that AI applications are capable of discerning patterns in the dataset that humans would not be able to discover."

the AI application works and on the basis of what criteria a decision is arrived at.[8]

The data subject also has the right not to be subjected to a decision based exclusively on automated processing which has legal effects for him or which otherwise affects him substantially.[9] Although this article is formulated as a 'right', according to the European privacy regulators, there is a de facto general prohibition on fully automated decision-making, notwithstanding exceptions.[10] The termination of an employment contract exclusively on the basis of an automated decision is therefore prohibited. Targeted advertising based on profiling will generally not have a significant impact on the data subject and consequently does not generally fall under a prohibited form of automated individual decision-making.[11]

Automated decisions are usually reached on the basis of pre-programmed decision rules. In that case, it is relatively easy to inform users about the 'underlying logic' of the system and to monitor whether the system has arrived at a qualitatively good decision. It can be more difficult to satisfy these obligations if an AI application, trained on the basis of data, has taken the decision.

In that case, the relationship between the input and the ultimate output (the decision) is more difficult to trace – including for the developer of the application. After all, the decision-making becomes a kind of black box.[12] For example, an AI system developed by an insurer decided that drivers of red cars should have to pay a higher premium than drivers of cars of other colours.[13] This brings me to the second legal framework.

Privacy as fundamental right

The European Court of Human Rights (ECtHR) interprets the right to privacy as contained in Article 8 of the European Convention for the Protection of Human Rights and Fundamental Freedoms (ECHR) broadly. Based on that, citizens not only have a certain right to be left alone by the government, but they also derive from that a right to personal development. The ECtHR also tends to draw other fundamental rights within the scope of Article 8 ECHR, such as the right not to be discriminated against and the freedom of speech.[14] Finally, the right to privacy is not only relevant in constitutional and administrative law (vis-à-vis the government), but the right can also inform private-law legal relationships (between citizens and businesses), the so-called horizontal effect of fundamental rights.

We have now discovered that many datasets are not neutral, but reflect societal prejudices and socio-economic inequalities. An AI application quickly adopts this predisposition, possibly in reinforced form. Amazon, for instance, used an application to assess CVs but it emerged that the self-learning program had a conspicuous preference for men.[15] This problem is known as 'garbage in, garbage out' or 'bias in, bias out'.

The European privacy regulators expect controllers that use AI applications to frequently audit their algorithms and datasets for possible prejudices and other discriminatory elements that could be contained in them.

Finally, AI technologies, like many other digital technologies, often use insights derived from neuroscience and psychology with the aim of getting users to spend time and money in a virtual environment. Examples are the automatic modification of content based on profiles and the Fear Of Missing Out (FOMO) principle. It is not inconceivable that the use of such technologies could, under certain circumstances, be in violation of the right to privacy as interpreted by the ECtHR. The alternative is also conceivable: that, on the basis of Article 8 ECHR, a person could claim a right to access to a friendly and helpful care robot.[16]

Conclusion

Self-learning AI applications thrive best in data-rich environments. It is therefore no surprise that this category of artificial intelligence raises countless privacy-law questions, especially since the regulations on data protection were themselves introduced relatively recently. To some extent, the problems are the same as the problems of any big data collection. Where AI

applications are to take decisions themselves, additional questions arise, for instance about how the duty to inform is fulfilled and how the rules for automated individual decision-making are applied. Depending on the application and the circumstances of the case, issues relating to fundamental rights (privacy) can also arise.

"Additional questions arise, for instance about how the duty to inform is fulfilled."

Points for attention:

- Self-learning AI applications process large volumes of personal and other data so the GDPR must remain a constant point of attention in the development of these applications. There are no simple solutions; in a general sense, the advice is therefore to always make a careful, expert and documented consideration.
- The GDPR stipulates specific obligations for the use of automated individual decision-making. It is relatively more difficult to satisfy these obligations with self-learning AI applications than with the use of traditional applications based on decision rules; however, it is not impossible.
- The fundamental right to protection of one's private life is interpreted broadly by the European Court; consequently, this right can also be relevant for the application of artificial intelligence. For example, if an AI application would result in discrimination or exclusion, if personal autonomy were impaired as a result, or if an AI application would have such positive effects for personal development that a claim to a right to access these applications could arise.

1 For instance, the California Consumer Privacy Act (CCPA) which came into effect on the 1st of January 2020 is partly based on the GDPR.
2 With regard to personal data and anonymisation see Article 29-WG, Opinion 4/2007 on the concept of personal data, 20 June 2007, and Article 29-WG, Opinion 5/2014 on Anonymisation Techniques, 10 April 2014.
3 For example, see: Licht op de digitale schaduw: verantwoord innoveren met big data [Light on the digital shadow: innovating responsibly with big data], Report from the expert group on Big data and privacy to the Minister of Economic Affairs, August 2016, page 21.
4 The use of these and other technologies are recommended by the Council of Europe, Guidelines on Artificial Intelligence and Data Protection, 25 January 2019, and the International Working Group on Data Protection, Working Paper on Privacy and Artificial Intelligence, 64th meeting, 29-30 November 2018.
5 Carlini, N. (2019). Evaluating and Testing Unintended Memorization in Neural Networks, https://bair.berkeley.edu/blog/2019/08/13/memorization/.
6 van Eck, M. (2018). Geautomatiseerde ketenbesluiten & rechtsbescherming: Een onderzoek naar de praktijk van geautomatiseerde ketenbesluiten over een financieel belang in relatie tot rechtsbescherming [Automated chain decisions & legal protection: A study into the practice of automated chain decisions on a financial interest in relation to legal protection], page 30.
7 Article 13(2)(f) GDPR and Article 14(2)(g) GDPR.
8 Article 29-WG, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, 6 February 2018, page 25.
9 This right is contained in Article 22 GDPR. Incidentally, the right does not apply to an unlimited extent. The right does not exist if i) the decision is necessary for the establishment or performance of an agreement between the data subject and a controller; ii) the decision is permitted pursuant to a statutory provision that applies to the controller and that also provides for appropriate measures protecting the rights and freedoms and justified interests of the data subject; iii) the decision is based on the explicit permission of the data subject; or iv) the decision, other than on the basis of profiling, is necessary in order to satisfy a statutory obligation borne by the controller or one that is necessary for the fulfilment of a task carried out in the public interest. In these cases, several additional conditions may apply, such as that the data subject has the right to make his/her position known and that the automated decision cannot be based on any special categories of personal data.
10 Article 29-WG, Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, 6 February 2018, page 19.
11 Ditto, page 22. The situation could be different under certain circumstances, however. For example, if prices are adjusted in a targeted manner, making a particular product or service unaffordable for an individual.
12 Castelvecchi, D. (2016). Can we open the black box of AI? Nature, 538, pages 20–23.
13 Interview with E. Haasdijk, National AI course, track 6.
14 Van der Sloot, B. (2015). Privacy as Personality Right: Why the ECtHR's Focus on Ulterior Interests Might Prove Indispensable in the Age of "Big Data". Utrecht Journal of International and European Law, 31(80), pages 25–50.
15 www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G.
16 This topic is explored in Van Est, R. & Gerritsen, J.B.A. (2017). Human rights in the robot age: Challenges arising from the use of robotics, artificial intelligence, and virtual and augmented reality – Expert report written for the Committee on Culture, Science, Education and Media of the Parliamentary Assembly of the Council of Europe (PACE), The Hague: Rathenau Institute, pages 18-26.

DE CLERCQ

About the author
Jeroen van Helden serves as an associate at De Clercq Advocaten Notariaat. With a specialization in IT, Jeroen assumes the role of lead counsel in IT transactions and disputes, bringing to the table a wealth of expertise in software licensing, data protection, and compliance matters.