Kenniscentrum
Data & Maatschappij

ARTIFICIAL INTELLIGENCE AND DATA PROTECTION: AN EXPLORATORY GUIDE

Knowledge Centre Data & Society

October 2021

Citation: Knowledge Centre Data & Society, "Artificial intelligence and data protection: an exploratory guide", October 2021

Contact: Thomas.gils@kuleuven.be or jan.debruyne@kuleuven.be

www.data-en-maatschappij.ai

Table of contents

1. Introduction and structure of the guide

Context of the guide - There are almost daily reports about new systems and applications that use artificial intelligence (AI). This rapid development of AI systems is a good thing, bearing in mind the many benefits they may bring. Nevertheless, there are also a number of legal, ethical and societal challenges that need to be addressed. It is essential that AI systems are developed and used within the existing regulatory framework.

Because AI systems typically use large amounts of data, the General Data Protection Regulation (GDPR) is paramount. The GDPR protects fundamental rights and freedoms of natural persons and in particular their right to the protection of personal data. Provisions of the GDPR must therefore be complied with in the design, development and use of AI systems.

Various international data protection authorities have published studies and official policy documents on this subject in recent months. With this exploratory guide on AI and data protection, the Flemish Knowledge Centre for Data and Society (KCDS) also aims to clarify the application of the GDPR to AI systems.

Development of the guide - This guide was developed through consultation with and input from stakeholders and with the support of the Flemish Department of Economy, Science & Innovation (EWI). Following internal consultation, a table of contents was drawn up and a proposal of topics to be covered was made. This was distributed to the stakeholders. Based on their feedback, the table of contents and topics covered were changed and/or refined. Researchers at the KU Leuven Centre for IT & IP Law (CiTiP) are responsible for the content and coordination of this guide. Additional feedback, questions and input on/about this guide can be sent to them at any time. The sheets and practical tools based on this guide are distributed in consultation with researchers at VUB-SMIT. Finally, we would also like to thank David Stevens, Chairman of the Belgian Data Protection Authority, for his valuable feedback.

Objectives of the guide - This guide has two objectives. Firstly, it aims to provide organisations and users with information on applying the GDPR in the design, development and use of AI systems. Secondly, the guide constitutes the framework from which other practical instruments will follow. These practical information sheets in Dutch are available online.

Structure of the guide - This guide is further structured into four sections:

- Chapter 2 discusses the concept of artificial intelligence and several other fundamental concepts.
- Chapter 3 examines the scope of the GDPR and applies it, where appropriate, in the context of AI.
- Chapter 4 explores how to ensure data protection in the design and development of AI systems.
- Chapter 5 studies how data protection can be ensured when AI systems are used.

Where useful, each section starts with an overview box, explaining both the essence of the section discussed and a number of concrete actions. The applicable provisions of the GDPR are then discussed in detail. This *multi-layered approach* ensures that the guide, besides being a comprehensive (legal) analysis, also endeavours to be a practical instrument. In concrete terms, this means that it first needs to be checked whether a sheet has already been published by the KCDS on a given subject. If this is not (yet) the case, then the practical steps in this report can be looked at. Additional information can then be found in the respective sections. The general bibliography indicates for each chapter which official policy and government documents were taken into account. Footnotes were used to refer to the relevant provisions in the GDPR or to specific authors/sources (other than the rather general policy and government documents).

Limitations of the guide - It is not possible to cover all topics on data protection and AI (such as the role and tasks of the data protection officer or binding corporate rules) in this guide. It was decided to address topics of specific interest to AI, whereby general guides and tools published elsewhere can be consulted in relation to general GDPR-related questions. Topics not covered may still be addressed separately via sheets and/or other practical tools.

About the KCDS - The Knowledge Centre for Data and Society is a collaboration between three university research groups: imec-SMIT-VUB, KU Leuven CiTiP and imec-MICT-UGent. It is part of the Flemish Policy Plan on Artificial Intelligence, and receives support from the Flemish government (EWI). The KCDS is the central hub for the legal, societal and ethical aspects of data-driven applications and AI applications.

2. What is Artificial Intelligence?

According to the European Commission (EC), artificial intelligence refers to systems that display intelligent behaviour by analysing their environment and taking actions – with some degree of autonomy – to achieve specific goals. AI-based systems can be purely software-based, acting in the virtual world (e.g. voice assistants, image analysis software, search engines, speech and facial recognition systems). Or AI can be embedded in hardware devices (e.g. advanced robots, autonomous cars, drones or Internet of Things applications).

The High-Level Expert Group on Artificial Intelligence[1] applies a **broader definition**. Artificial intelligence (AI) systems are software (and possibly also hardware) systems designed by humans that, given a complex goal, act in the physical or digital dimension by perceiving their environment through data acquisition, interpreting the collected structured or unstructured data, reasoning on the knowledge, or processing the information, derived from this data and deciding the best action(s) to take to achieve the given goal. AI systems can either use symbolic rules or learn a numeric model, and they can also adapt their behaviour by analysing how the environment is affected by their previous actions.

An AI system can be **strong or weak**.[2] A strong AI system is intended to be a system that can perform most activities that humans are able to do. Strong AI systems do not exist yet. Weak/narrow AI systems are instead systems that can perform one or few specific tasks. Examples are self-driving cars or facial recognition applications.

AI as a scientific discipline has **several subfields** including natural language processing[3], expert systems[4] and robotics.[5] Roughly speaking, a distinction can be made between a '**knowledge-based**' and a '**data-based**' approach to AI. The first approach attempts to map the knowledge of a human expert as best as possible through observations and conversations with the expert, and then attempts to cast that knowledge into representations, rules, and search strategies that approximate the expert's behaviour. The second data-based approach in particular, machine learning (ML)[6], receives a lot of attention these days.[7] This technique starts from data[8] about people's behaviour, about the decisions they have made or about phenomena observed via sensors. Statistical techniques are then used to identify patterns in the data, and these patterns are then used to solve new problems.[9]

ML is a subcategory or type of artificial intelligence. ML is based on algorithms that are able to learn based on previous experience, the so-called self-learning algorithms. This allows computers to learn without being explicitly programmed to do so.[10] The more data these systems or tools process, the better the algorithms in these systems will detect patterns in the collected data. They do this autonomously, without instructions, but with the help of examples or suggestions. The key elements, according to the recent White Paper on AI, which outlines the EC's strategy for AI, are **data** and **algorithms**. Data is any form of information that can be processed by a computer. This can range from a few minimal datasets to multiple datasets. Processing a huge amount of data is

[1] Better known as the *AI HLEG*.

[2] Sometimes the terms 'generalist' and 'specific' AI systems or even 'narrow' or 'general' AI systems are also used.

[3] The ability to process and produce spoken and written language.

[4] Systems that have knowledge of a given field and can apply that knowledge to the facts of a case by reasoning, for example in a medical context.

[5] M.J. Vetzo, J.H. Gerards and R. Nehmelman, *Algoritmes en grondrechten*, Boom Juridisch, The Hague, 2018, p. 43.

[6] Machine learning is often abbreviated to ML.

[7] According to imec, AI (translation) "refers to machines that can learn, reason, make decisions and act on their own thanks to an exceptional understanding of data", i.e. without someone having to tell them what to do each time. Note that the word "data" appears in this definition. Without data, there is no AI". See: https://www.imec.be/nl/artikelen/wat-is-artificiele-intelligentie-en-wat-ben-ik-ermee.

[8] Data set and dataset are used as synonyms in this guide.

[9] L. Steels, "Artificiële intelligentie. Naar een vierde industriële revolutie?", Royal Flemish Academy of Belgium for Sciences and Arts, 2017, p. 14-17.

[10] M.J. Vetzo, J.H. Gerards and R. Nehmelman, *Algoritmes en grondrechten*, Boom Juridisch, The Hague, 2018, p. 43.

known as **big data**. An algorithm is a sequence of rules and instructions that achieve a predetermined goal. An algorithm reads, searches and sorts data to create knowledge.[11]

Deep learning (DL) systems[12] are currently the most common and advanced form of ML. DL uses 'artificial neural networks'. These are networks of digital neurons, inspired by the human brain. In this regard, the deep learning algorithm performs a layered analysis, using results from one layer as input for the analysis by a subsequent layer. This makes it possible to identify complex, hidden relationships in large datasets.[13] For example, a network that needs to recognise a traffic sign will focus on shapes, colours, and sizes. A first layer might look for an inverted triangle, a second for bright red, and a third for white. Each layer will indicate whether it has found its specific item, and how sure it is of it. For example, a neural network for image recognition can look at a photograph and recognise a right of way sign.[14]

As such, AI systems generally use **large amounts of data** that give these systems the ability to learn and become intelligent. This does not necessarily involve personal data, for example meteorological or financial data that are not linked to individuals. But, if the AI system uses **personal data**, as already stated, the GDPR applies. The GDPR must therefore be complied with in the design, development, roll-out and use of AI systems.

APPLICATIONS IN THIS GUIDE

There are numerous applications of AI systems to which the data protection provisions discussed in this guide apply. It is not possible to discuss all these AI systems in detail. It is important to bear the following questions in mind in the design, development and use of an AI system. Depending on the response, it can be determined whether and which requirements from the GDPR are relevant:

- Purpose: what is the application used for?
- Channels: what channels are used to collect data and to approach individuals?
- Which data: which data is processed by the AI system in the various phases?

This guide uses two case studies to clarify and illustrate, where appropriate, the application of the GDPR to AI systems.

AI in e-commerce

A first application is the use of an AI-driven e-commerce sales program. Familiar examples include Amazon Web shop, Alibaba or Collect&Go from Colruyt. There are also various online platforms that incorporate 'cognitive computing technology'. Cognitive computers learn from the data presented to them, both from structured sources (such as documentation, manuals, specifications) and unstructured sources (such as blogs, reviews, social media). They attempt to understand the context based on this (big) data. The main characteristics of cognitive systems are that they understand, learn, and reason and interact with humans in ways that are natural to us.[15] For example, the Expert Personal Shopper (XPS), a platform and internet bot uses online conversations with people to find out what they want to buy and then help them with it.

[11] For more information on these terms, see: M.J. Vetzo, J.H. Gerards and R. Nehmelman, *Algoritmes en grondrechten*, Boom Juridisch, The Hague, 2018, p. 244.
[12] Deep learning is often abbreviated to DL.
[13] M.J. Vetzo, J.H. Gerards and R. Nehmelman, "Algoritmes en grondrechten", Boom Juridisch, The Hague, 2018, p. 43.
[14] The example was quoted from (translation): https://www.techzine.be/blogs/trends/25516/ai-machine-learning-en-deep-learning-wat-is-het-verschil.
[15] IBM "How to get started with cognitive technology”, https://www.ibm.com/watson/advantage-reports/getting-started-cognitive-technology.html. See also: https://searchenterpriseai.techtarget.com/definition/cognitive-computing.

1. The goal of such an AI system is to enhance customer experience and increase sales by personalising elements such as:
    a. recommending suitable articles to visitors of a web shop;
    b. recommending suitable articles to go with other articles already in the shopping basket;
    c. providing targeted discounts to customers;
    d. Creating so-called lead generation[16], i.e. attracting potential customers by providing them with relevant content.

2. Various channels can be used in such AI-driven e-commerce sales programs, including:
    a. Web shops including the use of pop-ups, chatbots or product notifications;
    b. E-mail correspondence, primarily with existing customers;
    c. Websites and applications from third-party service providers such as Google Ads or social media such as Facebook, Twitter or Instagram.

3. Various personal data may be used and processed by an AI-driven e-commerce sales program including:
    a. previous purchases of the user;
    b. the user's behaviour on the website (e.g. clicking, returning, viewing certain items again, etc.);
    c. the user's online information (via Google for example);
    d. information entered by the user into the web application;
    e. location data;
    f. comparison with data from 'similar' users;
    g. comparison with similar combined purchases/views (e.g. other customers who bought/viewed X, also bought/viewed Y);
    h. biometric data, such as facial recognition applied to a profile picture.

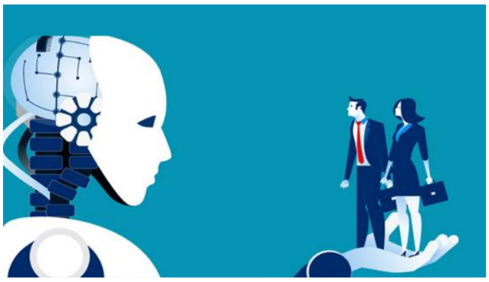

AI in recruitment

A second application is the use of AI-driven systems in recruitment. There are various applications in this regard. For example, Skeeled is AI-based recruitment software used in different parts of the recruitment process such as pre-screening, ranking applicants or providing feedback to recruiters. AI can also help clarify unclear job descriptions before they are published. For example, the VDAB has been using AI - Jobnet - since late 2018 to improve and optimise the results of their automatic matching system. Yet sometimes this can also go wrong on account of bias.[17] For instance, an algorithm developed by Amazon to scan cover letters was found to disadvantage women.

1. The goal of using an AI system is to recruit the right profiles and thus make the process more efficient by, for example:
    a. showing vacancies online to the relevant people;
    b. making a ranking and/or selection from the applications and then choosing a suitable candidate;
    c. giving feedback to the people involved in the recruitment process.

2. Various channels can be used to collect data from a candidate, including:

[16] Lead generation is the process of identifying potential customers and generating their interest in an organisation's products or services.
[17] Bias refers to unconscious reasoning errors/inherent prejudice built into the AI system.

a. direct channels such as an online application page or emails;
b. social media platforms (e.g. Facebook, Twitter, LinkedIn and Instagram).

3. Various personal data may be used and processed by AI-driven recruitment systems, including:
   a. Information provided by a candidate in a CV or cover letter, such as identity, hobbies, photos, previous work experience and education;
   b. Information from social media profiles such as the amount and content of posts, friends, profile, social and geographical environment or photos;
   c. Other personal online information such as an e-commerce profile;
   d. Other information sourced from third parties.

3. To which activities does the GDPR apply and which roles can an organisation fulfil under the GDPR?

Essence

It is crucial to know what kind of data is or will be used in a given AI system. If the system processes personal data, then the GDPR must be complied with.

The term 'processing' is very broad. Personal data is processed as soon as something is done with this data or even as soon as the data passes through an environment controlled by the organisation, even if there is no effective access and the organisation does not do anything else with the personal data.[18]

Personal data are both data that make it possible to identify a natural person and data that relate to an identified or identifiable person. Individuals can be identified by name or address ('direct identification'), but also by their IP address, cookie identifier or other factors ('indirect identification'). If a person cannot be immediately identified, it must be verified whether (indirect) identification is possible or not.

Pseudonymised data fall under the application of the GDPR. The GDPR does not apply to anonymous data. However, in such cases it must be verified that no re-identification is possible. The problem is that big data facilitates the possibility of re-identification through the combination of different datasets. As such, anonymising personal data is not always permanent and not every anonymisation method is an appropriate method for protecting data. Indeed, data circulates on the Internet, is traded, new datasets are created and third parties may be in possession of information that makes it possible to link data, of which the original controller is unaware. This means that it will become increasingly difficult to draw clear boundaries between personal and non-personal data.

AI systems can also use 'mixed datasets', containing both personal and non-personal data. This means for these types of datasets that: (i) non-personal data are covered by the Regulation on the free flow of non-personal data and (ii) personal data are covered by the GDPR. If both datasets are 'inextricably linked', the GDPR will apply to the entire dataset, even if personal data make up only a small part of the set. There is a good chance that an AI system will use mixed datasets that are 'inextricably' linked, and the GDPR will therefore apply.

One of the most important aspects under the GDPR is defining the different roles and responsibilities with regard to the processing of personal data. It is therefore crucial to determine whether an organisation is a controller or a processor, as this has a fundamental impact on the obligations that must be complied with.

In action

- ✓ Make an inventory of the data that the AI system will use.
- ✓ Verify which data is personal data and which is not.
- ✓ If it is personal data, determine whether all of this data is actually needed for the development and operation of the AI system.
- ✓ Verify whether the personal data used in the training phase can be anonymised or at least pseudonymised for the operational phase, if they are still used in these phases.
- ✓ Check if there is a way to anonymise/pseudonymise some/all data without a major (functional or technical) impact on the AI system.
- ✓ Analyse what probability of re-identification if anonymous data is used.
- ✓ Check whether mixed datasets can be separated or are 'inextricably' linked.

[18] Art. 4(2) GDPR.

3.1. What are personal data?

The GDPR applies to the **processing of personal data.**[19] The definition of processing is very broad[20] and consequently, the GDPR will apply to almost any operation involving personal data.

The term personal data means all **information about an identifiable living natural person**. Examples of personal data include:

- a first name and surname;
- a home or delivery address;
- an e-mail address such as firstname.surname@company.be;
- identity card number;
- location data (for example, the location data on a mobile phone);
- Internet protocol address (IP address);
- identification cookie[21];
- advertisement ID from telephone;
- data held by a care or service provider, for example in the form of a symbol that gives someone a unique identity;
- one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person,
- literally any other information that can be linked to a natural person.

It is not always straightforward to determine whether an element of data is personal data. In any case, it is advisable to **handle personal data with caution** and to ensure that there is a clear reason for processing.

Personal data consists of a **number of important building blocks**:[22]

ALL INFORMATION

First, it relates to all information, which highlights the fact that personal data is a broad concept. In terms of the **nature** of the information, this includes data of all types about a specific person. This information can be objective, such as a person's blood type, or subjective like opinions or judgements. As regards the **content** of the information, this can be data that provides information of any kind, even technical data. Furthermore, personal data can relate to an individual's private and family life, but also to the activities that an individual is engaged in, such as at their work, leisure time or as a consumer. The **form** of the medium on which the information is stored can be in any form such as alphabetical, numerical or graphic.

INFORMATION ABOUT A PERSON

Secondly, the information must relate to a natural person. It can be assumed that information relates to a person when it is **about that person**. In many situations, this relationship can be easily ascertained. For example, the data in an individual personnel file is clearly 'related' to the person's situation as an employee. In this regard, any information that makes it possible to identify a person directly or indirectly must be considered personal data.

Nevertheless, there are also situations when it is **not straightforward** to determine whether data relates to a person. Indeed, in some situations the information relates to objects such as a car or a house. These objects are usually owned by someone, are under someone's control, or exert influence over an individual. In such

[19] Art. 2 GDPR.

[20] Art. 4(2) of the GDPR describes processing as "any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction".

[21] For example, the use of cookies or similar technologies to track an individual across different websites involves the processing of personal data if such tracking is accompanied by online identifiers used to create a profile of the individual. For example, a social media 'handle' or an individual's username, which may seem anonymous or meaningless, is still enough to identify them, because it uniquely identifies that person. The user name is an element of personal data if it distinguishes one individual from another, regardless of whether it is possible to link the 'online' identity to a 'real world' person.

[22] The Article 29 Working Party is the independent European working group which until 25 May 2018 was responsible for handling questions relating to privacy and personal data protection. The Working Party was replaced by the European Data Protection Board. For more information see: https://edpb.europa.eu/our-work-tools/article-29-working-party_en).

cases, it can only be assumed indirectly that the information relates to a certain person, by linking it to other data that make identification possible, but here too it is personal data.

To determine whether data relates to a person, the following elements are important:

- **content** of the data (are they directly related to an individual or their activities?);
- **purpose** for which the data are processed;
- **outcomes** or **consequences** for the individual on account of the processing of the data.

IDENTIFIED OR IDENTIFIABLE LIVING PERSON

Thirdly, it must be an identified or identifiable living person. In general, an individual can be deemed to be identified when they can be **clearly distinguished** from other members of a group. An individual can be identified via **identifiers**. Examples are external characteristics or a quality of an individual that cannot be immediately ascertained, such as a name, position or profession. Identifiable implies the possibility that the person can be distinguished. Identification is also possible in an indirect way. This usually involves a small or large number of 'unique combinations'.

In cases where it is not possible at first sight to distinguish a given individual through the available means of identification, that person may nevertheless be identifiable because the **combination of that information** with other data (which may or may not be available to the controller) makes it possible to distinguish the data subject from other persons. A typical example is information related to objects. This is because objects are generally owned by, controlled by or exert an influence on, or have some physical or geographical proximity to, individuals or other objects. However, the information may lead to a person and in such cases can only be considered to indirectly identify those persons.

3.2. What are special categories of personal data?

There are a number of special categories of **data which in principle may not be processed** unless there are **exceptional grounds** for doing so. Other 'regular' data may in principle be processed, provided that the processing is done in accordance with the GDPR and there are legitimate grounds for processing, among other things.

These 'sensitive' categories of data are: personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, data concerning a natural person's sex life or sexual orientation, genetic data, biometric data for the purpose of uniquely identifying a natural person, or data concerning health.[23]

However, the prohibition on processing such data does not apply in a number of cases provided by the GDPR.[24] If an organisation wishes to process these personal data in these cases, it must in any case comply with all other principles of the GDPR, as well as any specific rules on the processing of such data.

In an **AI context**, the **following exceptions** may be relevant:

1.	If **explicit consent** was given by the data subject for the processing of those data for one or more specified purposes.
2.	If the data has been made **public** by the data subject.
3.	If the processing is necessary for **scientific research or statistical purposes** based on a statutory provision.
4.	If the processing of such data is necessary for the purposes of carrying out obligations and exercising **rights** of the controller or the data subject in the area of **employment law, social security law and social protection law.**

[23] Art. 9 and 10 GDPR.
[24] Art. 9 GDPR.

5. If the processing is necessary for purposes of **preventive or occupational medicine** or for the assessment of the **employee's fitness for work**.

Given that the processing of this type of personal data often entails an increased level of risk, it is likely that a **Data Protection Impact Assessment** (DPIA) will need to be performed.[25]

3.3. What is the significance of the distinction between anonymisation and pseudonymisation?

The distinction between anonymisation and pseudonymisation of data is important in the context of data protection. The GDPR uses the term **pseudonymisation** to refer to **encrypted data** that can **no longer be linked to a specific natural person** without additional information serving as a key. The additional information for attributing the personal data to a specific data subject is kept separately.[26]

APPLICATION

Suppose someone applies for a certain job position. The HR Department had an AI system developed that splits the application file into two folders. The first step is to delete the first page that contains the name and contact details, and keep the rest of the document in Folder 1. This document is given an automatically generated number in a second step, and is then forwarded to a recruiter. The HR Department will keep the first page of the application with name and contact details along with the automatically generated number in Folder 2. On its own, the information in Folder 1 does not make identification possible, but combined with the information in Folder 2, the applicant can be identified.

Pseudonymisation is therefore not a method of anonymisation, but **reduces the possibility of information being linked to the data subject.** Pseudonymised personal data for which a key exists to recover the original personal data remains personal data and is therefore subject to the obligations of the GDPR.

The GDPR attaches a **number of benefits** to pseudonymisation that may also be useful for AI systems**:**

1. There is a wider scope to process the data for a purpose other than that for which it was collected.[27]

2. This technique can help as a technical and organisational measure to implement the principle of purpose limitation and the obligations of data protection by design and by default.[28]

3. It helps to meet the requirements of data security.[29]

4. It is an important safeguard in the context of processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes.[30]

Anonymisation is not defined by the GDPR, but essentially means that the **natural person to whom the data relates is not or no longer identifiable**. The data are only truly anonymised when the anonymisation is irreversible. If an organisation is indeed capable of irreversible anonymisation, then the GDPR does not apply (anymore).[31]

[25] . See also section 4.4 on the DPIA.
[26] Recital 29 GDPR.
[27] Art. 6(4) GDPR.
[28] Art. 25 GDPR. See also section 4.1 on data protection by design.
[29] Art. 32, 33 and 34 GDPR. See also section 4.3 on security of processing of personal data.
[30] Art. 89(1) GDPR. See also section 4.5 on the processing of personal data for scientific or statistical purposes.
[31] Recital 26 GDPR.

Identifiability is therefore the criterion to assess whether data are pseudonymous or anonymous. In this regard, all objective factors, such as the cost and time required for identification, must be taken into account. The technology available at the time of processing and technological developments should be taken into consideration.[32]

Despite anonymisation, there may **still be a possibility of re-identification**. This is the process of reconverting anonymised data into personal data by using data matching or similar techniques. These techniques often use a form of ML, meaning that for some of these applications there is a risk of re-identification. For example, in 2019, researchers developed a model with which they could correctly re-identify 99.98% of Americans in any dataset, using 15 demographic characteristics.[33]

In this regard, the WP29 states that there are **three criteria** to be considered in determining whether re-identification can take place, namely:

- **Singling out**: the possibility to isolate some or all records which identify an individual in the dataset;
- **Linkability**: the ability to link, at least, two records concerning the same data subject or a group of data subjects;
- **Inference**: the possibility to deduce, with significant probability, the value of an attribute from the values of a set of other attributes.

According to the WP29, an anonymisation solution that rules out these three risks is sufficiently resistant to the risk of re-identification. However, the WP29 has itself indicated that reaching this threshold is very difficult. Indeed, each method entails at least a small risk of re-identification. Only a **combination of different techniques** would make it possible to completely anonymise personal data.

WHAT ABOUT BIG DATA?	One problem is that big data (especially when combined with the computational power of AI systems) increases the possibility of re-identification due to the possible **combination of different datasets**. As such, anonymising personal data is not always permanent and is perhaps no longer an appropriate method for protecting data.[34] Indeed, data circulates on the Internet, is traded, is integrated into new datasets, etc. Moreover, third parties may be in possession of information that makes it possible to link data, and of which the original controller is unaware, and which leaves open the possibility of liability. This means that it will become **increasingly difficult** to draw clear **boundaries** between **personal data** and **non-personal data**.

Consequently, the question that arises is whether in certain circumstances it might not be better to **pseudonymise data** rather than to strive for almost impossible anonymisation of data. Indeed, according to some, **anonymisation** potentially has **several drawbacks**:

- reduced ability to link data back to individuals, so it is not always known whether the data was collected in accordance with the GDPR in the first place;
- less insight into the origin of the data, transformations and transfers over time, making it more difficult to develop a responsible data policy;
- Due to the combinations of different datasets, there is a risk of re-identification, which can entail increased liability;
- The data may be less accurate depending on the anonymisation techniques used. As a result, the dataset can be less useful.[35]

[32] Recital 26 GDPR.

[33] L. Rocher, J.M. Hendrickx and Y. de Montjoye, “Estimating the success of re-identifications in incomplete datasets using generative models”, *Nature Communication*, 2019, vol. 10, nr. 3069, https://doi.org/10.1038/s41467-019-10933-3.

[34] See in this context: G. LaFever, “Anonymisation does not work for big data due to lack of protection for direct & indirect identifiers and easy re-identification vs pseudonymization”, gdpr.report/news/2019/08/12/anonymisation-does-not-work-for-big-data-due-to-lack-of-protection-for-direct-indirect-identifiers-and-easy-re-identification-vs-pseudonymisation.

[35]For the practical application of anonymisation and the techniques that can be used in this regard, see section 4.2. on data minimisation.

3.4. What about non-personal data and/or mixed datasets?

As mentioned above, AI systems do not necessarily use personal data. They can also use non-personal data. Non-personal data are defined as **data other than personal data**.[36]

These are data that do not relate (or no longer relate) to an identified or identifiable natural person, such as data on weather conditions, insofar as they are not linked to an individual of course. It may therefore also be personal data that have been anonymised. The specific and unique circumstances of each individual case must be taken into account in assessing whether the data have been appropriately anonymised.

In reality, it can sometimes be **difficult** to draw a **clear line** between personal data (and therefore application of the GDPR) and non-personal data (and therefore application of the rules on non-personal data). AI systems can also make use of '**mixed datasets**'.

APPLICATION

Some web shops may use services provided by third parties in the context of customer relationship management (CRM). In this regard, AI can be used to improve the output and effectiveness of CRM tools.[37] The data of a customer must be made available in the CRM environment. The data required for the CRM department includes all information necessary to effectively manage customer interactions. These include, for example, their postal address, e-mail address, telephone number, etc. However, it may also pertain to products and services they purchase, as well as sales reports including aggregate data, which is non-personal data. As such, the data in the CRM environment may include both personal data and non-personal data.

In this context, the European Commission has published some guidelines on the interaction between datasets consisting of both personal and non-personal data.

When a dataset **consists of both personal and non-personal data**, this means that:

- non-personal data are covered by the Regulation on the free flow of non-personal data;
- personal data are covered by the GDPR.

If both datasets are '**inextricably linked**', the GDPR will apply to the entire dataset, even if personal data make up only a small part of the set.

The term 'inextricably linked' is not defined. It may refer to a situation where a dataset contains both personal and non-personal data and the separation of these data is either:

- impossible;
- economically inefficient; or
- deemed technically unfeasible by the controller.

There is a good chance that **AI systems use mixed datasets that are 'inextricably' linked**, meaning that the GDPR will apply.

APPLICATION

For example, the information 'Master's degree obtained' is in itself a non-personal data when it is no longer clear from the dataset to whom this data referred. However, in an AI system for recruitment purposes, it is

[36] Art. 3(1) of Regulation 2018/1807 of the European Parliament and of the Council of 14 November 2018, on a framework for the free flow of non-personal data in the European Union, PE/53/2018/REV/1, OJ L 303, 28.11.2018, p. 59-68.
[37] See, for example: Commercient, "Why Artificial Intelligence Integration in CRM is the Future for your Businesss," https://www.commercient.com/why-artificial-intelligence-integration-in-crm-is-the-future-for-your-business.

necessary that the 'degree obtained' and the 'applicant' are linked to each other. Separating these is therefore not always possible, meaning that they are 'inextricably' linked, and the GDPR therefore applies.

3.5. What are the different roles an organisation can play under the GDPR in an AI context?

One of the most important aspects under the GDPR is defining the different roles and responsibilities with regard to the processing of personal data. The distinction between (**data**) **controller** and (**data**) **processor** is important, because they each have different obligations under the GDPR. For example, only if this capacity is confirmed do we know who or which organisation is responsible for the transparency or accountability obligations.

At the different stages of the life cycle of an AI system, the **controller** is the natural or legal person, public authority or other organisation that decides on the purposes and means of processing personal data.

A controller may outsource various tasks to third parties who will perform those tasks on behalf of the controller and in accordance with the controller's instructions. These third parties are the **processors**. If they further process these personal data or perform additional processing operations on their own behalf, they become the controllers for those processing operations.

Below, we briefly explain per phase/activity in the life cycle of an AI system who should be considered the controller and who should be considered the processor. The table highlights several common cases and can serve as a guideline for new situations.

PHASE / ACTIVITY	WHO IS THE CONTROLLER?	WHO IS THE PROCESSOR?
DEVELOPMENT/TRAINING /VALIDATION	The organisation that (further) develops, trains or validates the AI system and decides what personal data will be used to train the system (and therefore determines the purpose and means). If this organisation obtains a set of personal data from a third party, it will also have the status of controller when processing such data. If the development, training, validation or (further) development is outsourced to a third party organisation and this third party organisation decides which type of personal data is used in this regard, it becomes a controller.	The organisation to which the development, training, validation or (further) development is outsourced, provided that the client to whom such services are provided: (i) identifies the purpose of the processing activity and; (ii) determines the significant characteristics of the personal data to be processed. This is regardless of whether this client/controller transfers the personal data to the processor or the processor obtains it through its own channels and; (iii) the processor processes such data only for the purposes specified by the controller.
LAUNCH/RELEASE/ COMMISSIONING	Any organisation that integrates an AI system into its product or service and thereby processes personal data for its own purposes. If the AI system (whether or not part of a wider product or service) is sold or licensed and already contains personal data, both organisations exchange personal data and are both controllers.	Any organisation that makes an AI system available to a controller whereby the AI system is integrated into the latter's product or service, or any organisation that does so because it is necessary for the proper performance of its service, but that does not itself process personal data obtained from the controller for its own purposes.

	Even if, for instance, a licensor makes a system available to a licensee and only the licensee is the controller (see on the right), the licensor still also becomes a controller when it processes personal data obtained from the licensee for its own purposes (e.g. to measure the efficiency of the AI system).	An organisation (service provider) that makes an AI system available to another organisation (user) is neither a processor nor a controller if: (i) this system is installed locally and stand-alone at the user's premises; (ii) the service provider does not have access to the local installation, e.g. for maintenance.
PROFILING	The organisation that decides to process personal data through an AI system for its own purposes. The exception for purely personal or domestic activities does not apply to organisations (e.g., Amazon) that provide means (Amazon Echo/Alexa) to process personal data for their own purposes, where these means are typically used in the context of such personal or domestic activities, such as, for example, voice assistants.	See above.
AUTOMATED DECISION-MAKING	The entity that performs the automated decisions with respect to data subjects for its own purposes.	See above.

Moreover, if two organisations jointly determine the purposes and means of the processing through an AI system, they may be considered as **joint controllers**. This may be the case, for example, where an organisation cooperates with another organisation in developing a product or service for which both parties provide personal data for the training and/or validation of the tool, and where they jointly determine the purpose of such processing and combine their technical resources, without one party processing personal data solely on the instructions of the other.

In principle, only the controller can decide to opt (or not) for **a technical solution based on AI** (or any other technology) **in the context of a processing activity**. The controller is therefore obliged to act with due diligence when choosing the actual IT tool, and in particular when outsourcing the processing or acquiring the tool. Nevertheless, in certain circumstances a processor may decide on the technical means used. In such cases, the processor will also assume part of this responsibility.

A controller must therefore list and assess **any quality specifications of the relevant solution beforehand and determine** the (necessary) **scope of the processing**. Indeed, a controller must assume the consequences of the related decisions. With regard to the data subjects, it will not be able to evade its responsibility by claiming that it did not have the correct information or technical knowledge. It is the controller's responsibility to carry out an audit, if necessary, and to decide whether the intended system is suitable for the intended purpose, and does not, for example, disproportionately process personal data.

In any case, controllers and processors can **never shift their responsibility to the AI system itself** and thus cannot, for example, hide behind the possible complexity or inscrutability of an AI system to justify infringements of the GDPR.

4. How can data protection be ensured in the design and development of AI systems?

The following sections discuss a number of GDPR requirements that are important in the design and development phase of AI systems, namely data protection by design, data minimisation, data security, DPIAs and the processing of personal data for scientific research or statistical purposes.

4.1. What does data protection by design mean and how can it be implemented in AI systems?

Essence

The development and use of AI systems and the relevant processes must be designed to inherently provide as much protection to personal data as possible. This is the principle of data protection by design.

With every process and development, consideration must be given, from the design phase onwards, to whether and how they (can) affect the way personal data are processed. Based on this, the necessary security features are then built into the process or product.

Likewise, products and services must be designed to use the most data protection-friendly settings for standard use, so that end-users can only derogate from them if that is their explicit intention. This is the principle of data protection by default, which is part of data protection by design.

Applying data protection by design avoids a so-called 'tech debt'[38]: the cost of complying with the GDPR is taken into account from the outset. Moreover, it avoids the need to subsequently rewrite systems to comply with the GDPR at an (often) greater cost, if such rewrite is possible at all.

Action points

General[39]

- ✓ Make sure that all employees are aware of the importance of data protection, pay attention to the risks and take responsibility in this regard (training and awareness-raising).
- ✓ Provide practical internal documents and guidelines to be applied by staff when working with personal data.
- ✓ Put in place an effective data policy that limits access to raw data and allows for the regulation, recognition and tracking of data, its access and its use.[40]
- ✓ Anonymise and pseudonymise personal data whenever possible.[41]
- ✓ Put in place an effective IT policy including effective technical and organisational security, role-based policy, encryption and staff awareness.[42]
- ✓ Make use of state-of-the-art technological and security applications, and ensure that this remains the case.[43]
- ✓ Incorporate data protection as a product requirement in every new development and process.
- ✓ At each stage and with each development, ask whether its implementation at that time or at a later time may have an impact on data protection.

[38] So-called 'technical debt' or 'tech debt' is a term in software development that reflects the implicit cost of additional reworking caused by choosing an easy (limited) solution during development rather than a better approach that might take longer. By not making certain efforts at the outset, a certain debt is built up, since this will have to be rectified in the ultimate version.

[39] Data protection by design is primarily a process obligation which implies that other substantial obligations are included in the design of a product, service or process. These action points therefore always relate to other obligations and principles from the GDPR, such as data minimisation, the obligation to set up an adequate security policy or the obligation to document how the GDPR is complied with. These are discussed below.

[40] See also sections 4.2. on data minimisation and 4.3. on security of processing.

[41] See also section 5.2. on storage limitation.

[42] See also sections 4.2. on data minimisation and 4.3. on security of processing.

[43] See also sections 4.2. on data minimisation and 4.3. on security of processing.

- ✓ Provide clear guidelines to determine when a DPIA needs to be performed.[44]
- ✓ Implement safeguards and so-called *nudges*[45] in systems for end users so that they are made aware that certain actions may pose a risk to data protection.
- ✓ Document the data protection analyses, considerations and choices made at each stage to demonstrate that the data protection by design obligation has been fulfilled. This also makes it possible to find out why certain measures were taken or not. Justify your decision in each case.
- ✓ Set the default settings of user software so that personal data is processed by default in the most protective way and the end user must explicitly deviate from these in order to process personal data in a less careful way.
- ✓ Impose the obligation on vendors to ensure that their services and products comply with the requirements of the GDPR, including the obligation of data protection by design and by default.
- ✓ Make use of so-called Privacy Enhancing Technologies (PETs).[46]

Specifically for AI systems:

- ✓ Build AI systems in such a way that they make it possible:
  - to figure out (to a given extent) the rationale behind the generated results (transparency) and the personal data used in this regard;
  - for data subjects to exercise their rights;
  - to be trained with pseudonymised and/or encrypted data;
  - to be trained with as little data as possible, but high-quality data[47].
- ✓ Avoid (facilitating) re-identification when anonymous data are processed, or deriving other personal data from the available data, which must not be processed, such as sensitive data.[48]
- ✓ Ensure that the effectiveness of the AI system is sufficiently tested and false positives and negatives (*hidden failures*) are effectively ruled out.
- ✓ Know your data:
  - Monitor the origin of and rights to personal data used to train AI systems. Ensure that it is allowed to use the dataset.
  - Ensure that the persons collecting and handling the data are able to recognise personal data.
- ✓ Make sure you have clean datasets:
  - delete redundant data;
  - ensure sufficiently representative datasets;
  - check that datasets do not contain any bias or prejudice that might reinforce social inequality or discrimination.

- ✓ Provide the necessary documentation in the context of accountability:
  - Document the applied data protection analyses, the decisions and considerations made, and the DPIAs performed during development, testing, and maintenance.
  - Document the properties of datasets, the way in which they were cleaned and the reasoning behind this. Also keep at least one representative sample of the dataset, anonymised if possible.

[44] See also section 4.4 on the DPIA.

[45] *Nudges* are architectural choices in the software that encourage a specific behaviour. For example, checking a certain option in advance in the hope that the user will agree.

[46] See also sections 4.2. on data minimisation and 4.3. on security of processing.

[47] See also section 4.2 on data minimisation.

[48] See also section 3.2 on special categories of personal data

A. Building data protection into applications and processes

Data protection cannot just be a layer of varnish on top of an AI system, it needs to be an inherent part of it. In accordance with the obligation of data protection by design, the necessary measures must already be taken **when determining the means of processing**[49] in order to:

- ensure the necessary safeguards for GDPR-compliant processing of personal data; and
- incorporate the pursuit of data protection principles[50] into the planned processing.

This does not mean that it is only when an AI system is used that it can be examined how personal data can be processed in that system in an GDPR-compliant manner. **Compliant processing must be ingrained in the system** from the outset.

Even after the design phase, compliance with the data protection by design requirement must continue to be evaluated, taking into account any changing circumstances. Compliance with this requirement is therefore an **ongoing process and exercise** that runs throughout the entire life cycle of an AI system.

Producers of products, services and applications for other users who do not intend to use or implement these themselves, and therefore do not process personal data themselves, are strictly speaking not obliged to apply this obligation.[51] Software that processes personal data is therefore not, strictly speaking, subject to data protection by design requirements if the developer does not use this software himself. The users of this software however, who process personal data via this software, must comply with data protection rules. As such, from the user market, there will be a need to purchase products that comply with the requirement of data protection by design, even if, strictly speaking, the producers do not have to incorporate it. However, where the producer also provides services through its product and thereby becomes a processor, it must comply with this obligation.[52]

APPLICATION

In e-commerce, for example, software can be made capable of recognising 'sensitive' data and warning the user when sensitive data is likely being processed. In addition, GDPR-compliant notifications could be automatically added to advertising messages generated by the application the AI system is part of. It is also advisable that the person entering the data of prospects always confirms from which (pre-defined) source the data originates, via a selection menu, so that this person is aware of the origin of the data and the source can easily be traced.

As regards recruitment, software should ideally be capable of recognising 'sensitive' data and warning the user when sensitive data is likely being processed. Notifications can also be built into the user interface that warn the software user that certain conditions must be met, for example, when choosing to retrieve publicly available social media information from a candidate (which is not always allowed). When creating a job announcement, it is also possible to ensure that information is automatically added to inform candidates about the processing of their personal data, such as the reference to the applicable privacy policy.

B. Risk-based approach

The risk-based approach of the GDPR is strongly reflected in the data protection by design requirement, given that the scope of the corresponding obligations depends entirely on the context. The state of the art, the costs,

[49] The term 'at the time of determining the means of processing' refers to the time at which the processes, techniques and methods that will be used to process personal data are determined. This is normally done in the design phase.

[50] These principles are lawfulness, fairness and transparency, purpose limitation, data minimisation, accuracy, storage limitation, integrity and confidentiality and accountability (Art. 5 GDPR).

[51] See also recital 78 of the GDPR which states that these producers must be 'encouraged' to implement data protection by design.

[52] Recital 78 GDPR.

the purpose of the processing and the risks associated with the processing must be taken into account in this regard, among other things.[53]

The higher the risk, the more effort is required from a developer. An AI system that processes fewer personal data and/or does not interact directly with people will require less extensive adaptation than a system that will process personal data or sensitive personal data[54] intensively and/or interacts directly with people.

Cost also plays a role in assessing whether the requirement of data protection by design is met. Cost does not only relate to the financial aspect, but to all efforts made, including time spent and staff costs. It cannot be expected from a small company with limited resources that they are able to build in the same safeguards as a large international corporation. However, the requirement to take appropriate effective measures must be adhered to. Indeed, not having adequate resources is not a justification for not complying with the requirements of the GDPR.

APPLICATION

In e-commerce, an advertising campaign based on simple criteria such as the indicated place of residence of an existing customer requires less extensive security measures and precaution than a campaign that also uses, for example, location data or data obtained from third parties. For recruitment applications that make a (pre) selection from submitted applications, more security safeguards need to be built in than applications that only give a score and leave the actual selection of candidates to the user.

C. Data protection by design covers compliance with all provisions of the GDPR

For data protection by design, the GDPR explicitly refers to pseudonymisation[55] and the principle of data minimisation.[56] However, this does not mean that only these measures have to be taken into account to fulfil the requirement of data protection by design.[57] It is clear that a **robust data minimisation policy** on the one hand, and **pseudonymisation** on the other hand are essential in order to comply with this obligation.

Data protection by design is therefore not so much a substantive obligation, but rather a **process obligation** that requires all obligations arising from the GDPR such as transparency[58], lawfulness or security[59] to be built into processes as much as possible.

How these obligations should be implemented is **up to the discretion of the organisation**. Technical and organisational measures can include all possible actions, ranging from explaining to staff how to process customer data to the use of sophisticated technical automated solutions. The measures do not have to be sophisticated in this respect. The only requirement is that they are effective and adequate in ensuring the obligations of the GDPR.

However, the **state of the art must nonetheless be taken into account**.[60] This means that technological developments must be taken into consideration. The measures taken to ensure compliance with the GDPR are therefore dynamic and must be adapted as necessary to the state of the art. The obligation to take account of the state of the art applies to both technical and organisational measures.

[53] Art. 25(1) GDPR.

[54]See also section 3.2 on special categories of personal data

[55] Art. 25(1) GDPR. See also section 4.1 on data protection by design.

[56] Art. 25(2) GDPR. See also section 4.2 on data minimisation.

[57] Recital 28 GDPR. This is reflected, inter alia, in the broad general references to having to take "technical and organisational measures", ensuring "the" data protection principles and building in safeguards to protect "the" rights of individuals and to comply with "the" requirements of the GDPR.

[58] See also section 5.1. on transparency.

[59] See also section 4.3 on security of processing of personal data.

[60] Art. 25(1) GDPR.

APPLICATION

In both e-commerce and recruitment, data protection by design requires safeguards to be built in at every stage. This can be done on various levels:

- At the **technical level** in the software itself by: recognising certain data; storing data in encrypted form; logging access to personal data; applying a role assignment policy so that each user profile only sees the data required by the role; requiring secure passwords; providing warnings to users and nudges that encourage proper use (e.g., building in additional steps when accessing sensitive data); providing retention periods for the different types of personal data and alerting the user when these are reached.
- **At the organisational level**, e.g. by ensuring that staff using this software have sufficient knowledge of data protection to correctly apply the applicable rules, both when using the software and beyond; and by providing the necessary internal procedures in this regard.
- **By having the software correctly process data at each stage**: When creating and storing user accounts; when entering customer data, leads, candidates and interesting profiles; when planning advertising campaigns; when generating ads (by including mandatory statements in these ads); when choosing the selection criteria for the target audience; when displaying ads; when using data to determine a consumer profile or recruitment profile; and when providing information to the target audience about why they are seeing certain ads, how long their data is processed and how they can unsubscribe.

Specifically for recruitment, rejected candidates can be automatically informed if they will be included in a recruitment reserve (or other database) and for how long. If necessary, an automated message can be sent every two years to candidates in the recruitment reserve to ask them whether they wish to remain in it.

D. Risk assessment

The data protection requirement under the GDPR is risk-based. Consequently, it must be verified at every development, every stage and for every product or process, including for AI systems in their entirety, whether there is an **impact on the data protection rights** of those whose personal data will be processed. In this way, proper data protection becomes a **(product) requirement** of the product or process to be developed.

Applying so-called **threat modelling** is a method to analyse the potential impact of an AI system on data protection.[61] To this end, threat modelling methods such as the ***LINDDUN Privacy Threat Modelling*** method for software development can be used.[62] Depending on the outcome, it may then be decided not to take any measures, or only specific measures, to increase the level of data protection.[63]

If there is a likelihood that an envisaged processing would pose a high risk to data protection, it should also be analysed whether a **DPIA is necessary**.[64]

APPLICATION

In e-commerce and recruitment, it is important to:

[61] Threat modelling is a method by which potential risks such as structural weaknesses or the lack of appropriate protective measures can be identified, listed, evaluated and prioritised according to their risk.

[62] For more information: https://www.linddun.org. See also: K. Wuyts, *Privacy Threats in Software Architectures*, Ph.D., 2015.

[63] This may involve weighing up different elements and, for example, deciding not to take certain measures that would have a positive impact on data protection because the costs would be disproportionate or the impact minimal.

[64] See also section 4.4 on the DPIA.

- involve the compliance officer in product development processes from the outset and analyse the potential impact on data protection from the ideation/creative phase onwards;
- evaluate every year the processes and practices in place and their impact on data protection;
- evaluate new functionalities that are added to the software used (don't just assume they are data protection friendly);
- require suppliers to comply with the GDPR and ensure their products meet applicable requirements.

E. Document the evaluations made and measures taken

Under the general 'accountability' requirement of the GDPR, it **must be possible to prove** compliance with the data protection by design obligation.[65]

It must therefore be able to prove that **measures** have been taken and that they are **effective**. In this regard, it is necessary to include the following in reports:

- that the impact on data protection was analysed;
- which measures were taken;
- what considerations were made in this respect; and
- what the results are.

In this way, the **necessary information** can be subsequently **retrieved** on the one hand, and on the other hand it can be **demonstrated** how data protection was considered and why certain measures were or were not taken.

APPLICATION

In e-commerce and recruitment, it is therefore important to:

- ensure that a report is made of the considerations and risk analyses of the product development;
- retain notifications to candidates, customers or prospects; document actions such as the following: record both the subscription and unsubscription by a candidate or prospect to given communications; record this in a file with indication of the time and copy of (if any) related correspondence.

F. Avoid 'technical debt' by applying data protection by design

By taking into account the applicable GDPR principles from the development stage of AI systems, burdening the final product with so-called technical debt[66] can be avoided. If data protection by design is not taken into account, the system will have to be subsequently adapted to meet the requirements of the GDPR, which is often a complex or even impossible task. Moreover, the cost of such a ‘rewrite’ can be considerable.

G. Data protection by default

When AI systems are developed, it must be ensured that the **default configuration provides the most data protection-friendly settings**. This is the requirement of data protection by default.[67]

As such, when **AI systems are used with their default settings**, they must, for example:

- process personal data only within the limits of the intended lawful processing;
- collect and process only those personal data that are required for such processing;

[65] Art. 5(2) GDPR. The accountability principle is one of the basic principles of data processing. It requires controllers and processors to be able to demonstrate that they have taken steps to comply with the obligations under the GDPR.
[66] For more information on this concept see above footnote 38.
[67] Art. 25(3) GDPR.

- not retain personal data for longer than necessary.[68]
- limit access to (non-pseudonymised) personal data;
- not publicly disclose the personal data.

APPLICATION

In e-commerce for example, it is recommended, by default, to only send communication to individuals, when, at the moment when their data is entered into the system, an option field is filled in with a value that allows these individuals to be contacted (for example 'consent given' or 'existing customer'). It is therefore important that, by default, no communication can be sent to people who have objected, for example by clicking the unsubscribe link in an e-mail. In the case of recruitment, it is advisable to provide standard fields in job advertisements to refer to the privacy statement and text fields for a brief explanation of the processing of data.

H. Certification

Certification[69] can be used as an **element** to demonstrate that the data protection by design requirement has been fulfilled.[70] However, even in the case of certification, it must still be demonstrated that the privacy by design obligation has been effectively fulfilled. But the **burden of proof will be lighter**. To date, there are no standards for data protection by design, and it is not possible to have a product or service certified as being compliant in this regard.

4.2. What does the requirement of data minimisation regarding personal data mean for AI systems?

ESSENCE

Data minimisation must be applied when personal data are collected, used and stored. The internally granted access to these personal data must also be limited. For each of these aspects, personal data may in essence only be collected, used, stored and consulted insofar as this is necessary for the purposes for which they are (or may be) processed. As such, if a similar outcome can be achieved without using (given) personal data, these personal data may not be used for that purpose.[71]

The personal data processed and the access to it must be limited to what is strictly necessary. In this regard, the unnecessary duplication of personal data must be avoided. Personal data must not be kept longer than necessary,[72] in line with the 'storage limitation'.[73]

ACTION POINTS

- ✓ Know your data: its origin, the rights attached to it, its properties and attributes must be known in order to know whether it has been correctly collected and for what purposes it can be used.
- ✓ Specify in advance which types of personal data (qualitative) and also which volume of personal data (quantitative) are strictly necessary to train and/or use the AI system. In this regard, rely on a multidisciplinary team, including experts in the domain where the AI system will be applied. Justify and document.
- ✓ Evaluate whether techniques and/or systems can be applied that make it possible to work with fewer personal data, to work with encrypted personal data or to use methods such as Generative Adversarial Networks (GANs), federated learning or transfer learning.[74]

[68] See also section 5.2. on storage limitation.
[69] As provided for in Article 42 of the GDPR.
[70] Art. 25(3) GDPR.
[71] Recital 39 GDPR.
[72] Recital 39 GDPR.
[73] See also section 5.2. on storage limitation.
[74] See Section D below for more information on these and other techniques and methods.

- ✓ Set up an effective pseudonymisation policy. Ensure that personal data can only be used, consulted and processed in pseudonymised form when it is not strictly necessary for this to be done with the 'raw' personal data.
- ✓ Set up an effective anonymisation policy. Anonymise or delete personal data for which it is no longer necessary to identify the persons to whom they relate or for which it is no longer justified to process them.
- ✓ Implement an enforceable role and access policy in the organisation, whereby individuals and applications can only access raw personal data if this is actually necessary, and following additional identification or authorisation. Others will only have access to pseudonymised personal data. Ensure that access to personal data is only possible with an individual account, and is logged.
- ✓ Clean up datasets regularly and do not process data longer than necessary.[75]
- ✓ Document and record each step and evaluation, in accordance with the accountability requirement.
- ✓ It is therefore important that:
    - justification can be given as to why certain personal data are being processed;
    - personal data are stored centrally in one place without unnecessary copies;
    - a proper ICT access policy is put in place which takes into account the different roles of staff in the processing chain;
    - an effective anonymisation and pseudonymisation policy is set up;
    - wherever possible, as little personal data as possible is used;
    - techniques are used that reduce the need for personal data, the volume and the risks of exposure.

A. General

Personal data must be **adequate**, **relevant** and **limited** to **what is necessary** in relation to the purposes for which they are processed.[76] They should be processed only if the purpose of the processing cannot reasonably be achieved by other means.[77]

Data minimisation is closely related to the principle of storage limitation,[78] the requirement of 'accuracy' of personal data and the prohibition on (re)using personal data for other purposes ('purpose limitation'). Data minimisation must be viewed **broadly** and applies not only to the **collection** of personal data, but also to the **internal access** and **use** of personal data. Data minimisation is an essential **component** of any data protection by design strategy.[79]

Data minimisation has the following **advantages:**

- Protection of the rights of the data subject:
    - there is no further intrusion into their privacy than what is necessary;
    - as few people as possible within an organisation have insight into their privacy;
    - who had access to what personal data is traceable.
- Reducing the risk that data breaches occur and the risk arising from data breaches:
    - information that is not collected cannot be breached;
    - personal data that is no longer relevant cannot be breached;
    - the likelihood that data breaches occur is reduced. This applies to insider threats (threats from within an organisation), external attacks (such as phishing and hacking) and unintentional data leaks;

[75] See also section 5.2. on storage limitation.
[76] Art. 5(1)(c) GDPR.
[77] Recital 39 GDPR.
[78] See also section 5.2. on storage limitation.
[79] See also section 4.1 on data protection by design.

- the risk of damage in the event of a data breach is reduced. For example, the damage is more limited when only pseudonymised personal data is leaked as a result of a phishing attack, because the person via whom the leak occurred did not have access to raw data.

B. Risk-based approach

The data minimisation requirement cannot be seen in isolation from the risk-based approach of the GDPR either.

Processing fewer personal data will in the first instance **reduce the risk** of **too much insight** into the **privacy** of the data subjects, and the risk that unexpected conclusions will be drawn from these personal data. Effective cleaning of datasets also contributes to the **quality of the data.**

This also reduces the **risk of a data breach** and any **damage** in the event of a data breach:

- Processing fewer personal data, anonymising and pseudonymising data more rapidly, encrypting data, securing and restricting access to personal data or storing personal data in one central location all help reduce the risk of a data breach.
- If a data breach nevertheless occurs, the number of persons affected, the usability of the leaked personal data and thus the extent of the damage to these individuals is significantly reduced.

Another consequence is that data minimisation must be implemented and ensured in a more far reaching manner as **the risks increase**. This is the case, for example, when:

- large volumes of data are processed;
- the data of a larger group of persons are processed;
- more types of personal data relating to the same person are processed;
- personal data are processed which belong to special categories or which, more generally, can be described as sensitive.

C. Minimum data protection and effectiveness of AI systems

It is often assumed that more data is (always) better for the functioning and results of AI systems.

However, this is not necessarily the case. By applying minimal data protection, training of the AI system with irrelevant parameters can be avoided. This avoids the risk that these parameters are regarded as significant by the AI system, and that it would incorrectly base its conclusions on them. Good knowledge of and cleaning up the dataset therefore contribute to the quality of the data and the results.

By carefully selecting which data will be processed, the so-called ***curse of dimensionality***[80] **can be avoided**. In this regard, an (AI) system is considered to have the best performance, when supplied with an optimal volume of data. If this optimal level is exceeded, efficiency is reduced.

The risk of ***overfitting*** can also **be avoided** by only working with the relevant parameters. By adding too many parameters, the AI system may be too adapted to the training data. It may then start to assign value to elements that are present in the training data, but which may be less relevant or not always present in other datasets. Consequently, such an AI system would be less effective when new datasets need to be analysed.[81]

D. Collect and use fewer personal data (training phase and deployment phase)

The **following questions and aspects** are relevant when collecting and using personal data.

[80] . More information: https://deepai.org/machine-learning-glossary-and-terms/curse-of-dimensionality.
[81] More information: https://www.datarobot.com/wiki/overfitting.

1. Can the processing of personal data be avoided altogether?

Collecting fewer personal data means processing fewer personal data. The first step in applying data minimisation is therefore to ensure that **as little personal data as possible** are processed.

The first question is therefore: do we actually need actual personal data? In other words, can the AI system be trained with other data? For example with synthetic personal data[82] or anonymised data[83]. Personal data should be processed only if the purpose of the processing could not reasonably be fulfilled by other means.[84]

Application

If an e-commerce AI system is trained on the existing customer data of a user, it is possible that it will not be trained on the actual data but on an anonymised dataset for this purpose.

Likewise in recruitment, an AI system can be trained locally on the personal data of the staff working in the organisation that wants to use the system. Here too, it is appropriate to anonymise this dataset in advance if possible.

2. Can the processing be done with fewer personal data?

If the use of 'real' personal data is actually required, the next question is whether the use of the personal data is **proportionate**. In other words, which of these personal data are actually necessary to achieve the intended purpose?

This question must be posed at **two levels**.

Firstly, the **number of persons** from whom data need to be processed must be looked at. Secondly, the question is also how much **different data** is required from each of these persons, i.e. how many *types* of personal data need to be processed. A **reasonableness test** also needs to be performed in this regard. It must be verified that the use of the envisaged personal data is proportionate in relation to the intended purpose and the risks involved for the individuals concerned.

As such, if the training can be done with personal data of x individuals, then no dataset can be used that contains the data of a larger number of individuals. If only y parameters are required from each person to achieve a qualitative result, no additional parameters can be used. Not even because this could reveal links that are not yet known. If the dataset to be used relates to more persons and/or contains more parameters, it will need to be **cleaned up** in order to be in compliance with the data minimisation principle.

Consequently, the amount and type of personal data required for the processing purpose must be analysed and specified **beforehand**. In this regard, it must be clearly established and explained why they are necessary, what needs to be learned from them and why these personal data are relevant and not excessive. Only personal data that pass this test can be processed.

In order to make this analysis, it is necessary to have **sufficient expertise** in the domain in which the AI system will provide analyses.

This test must be performed for **each individual processing operation**: just because it is justified to keep x parameters of y persons in a given dataset, it does not mean that this complete set can be used in other operations using the same dataset.

[82] *Synthetic data* are (in this case) data that mimic personal data, but are either produced, compiled from anonymous data or both.

[84] Recital 39 GDPR.

Application

In recruitment and e-commerce, the question of whether fewer personal data can be used is relevant in both training an AI system, and using an AI system. When using an AI system, the question is then how much personal data or parameters are actually needed to draw the desired conclusions. An effective application of reducing the personal data processed is to ensure that the input fields that a candidate or lead has to fill in only include those that are really necessary.

3. Do not retain personal data for longer than necessary

The personal data that can be processed after the above-mentioned tests cannot be processed for **any longer than necessary**. This is crucial in order to prevent personal data being accumulated for indefinite periods of time. If this were the case, it would create serious risks for data subjects, particularly in the event of a data breach.

Given the importance of this obligation, the GDPR has enshrined it in a separate principle, namely **storage limitation.**[85]

4. Destroy or anonymise personal data as soon as possible

When it is no longer necessary to link personal data to the data subjects, the data must be **destroyed**. This can be done by effectively **destroying (erasing) them completely** in all locations where they are stored. In this regard, measures can also be taken so that they are **not stored indefinitely in backup copies either**, and at the least, if this is impossible due to a legal obligation, made very difficult to access.

In practice, datasets will often be **anonymised**, which is a form of destruction from the standpoint of the GDPR.[86] Thanks to robust anonymisation, the personal data ceases to be personal data and the GDPR no longer applies in this respect.[87]

There are **various methods** that make it possible to anonymise personal data. Each has its advantages and disadvantages, which is why a combination of different methods is often used. A robust anonymisation policy requires that the techniques/methods used be adapted to the **state of the art** and that **re-identification tests be carried out**, to ensure the quality of the anonymisation. Given that a number of these techniques involve some form of aggregation and that any anonymisation results in some data being deleted, this will in many cases reduce the accuracy and comprehensiveness of the dataset.

Data sets can be anonymised using **open source applications** such as ARX[88] or Amnesia[89] or one of the many applications available on the market. An important consideration in choosing anonymisation software is that, as described above, different methods are applied side by side and that it is possible to perform re-identification tests. Of course, this is in addition to the primary requirement of achieving a quality end result.

There are various methods for measuring whether a dataset has been adequately anonymised, such as the concept of "k-anonymity".[90]

Below, **several methods** that are used individually or combined in order to anonymise personal data are explained, as well as some advantages and disadvantages of these methods.[91] Each of these techniques requires

[85] Art. 5(1)(e) GDPR. See below section 5.2. on storage limitation.

[86] See also section 3.3 on anonymisation and pseudonymisation.

[87] See also section 3.3 on anonymisation and pseudonymisation.

[88] See: https://arx.deidentifier.org.

[89] See: https://amnesia.openaire.eu.

[90]. See for example R. Shokri, C. Troncoso, C. Diaz, J. Freudiger and J-P Hubaux, “Unraveling an Old Cloak: k-anonymity for Location Privacy”, in: K. Frikken (ed.), *Proceedings of the 9th ACM workshop on Privacy in the electronic society (WPES 2010)*, 2010, p. 115-118.

[91] For more information, see for example the Personal Data Protection Commissioner Singapore, “Guide To Basic Data Anonymisation Techniques”, 25 January 2018, 39p.

that the **changes are definitive** and the anonymised data can no longer be linked to the original dataset. If not, it is not anonymisation, but pseudonymisation.

<table>
<tr><th colspan="2"><b><i>ATTRIBUTE SUPPRESSION, OR THE REMOVAL OF ALL NOTIFICATIONS OF A CERTAIN TYPE OF (IDENTIFYING) PROPERTIES</i></b></th></tr>
<tr><td><u>Concept</u>
This is the simplest form of anonymisation. The information deemed to result in identification is deleted.
<u>Example</u>
From a list of first name, last name and a test score, the last name is deleted. Only the first name and test score remain, meaning that (direct) identification is no longer possible.</td><td><u>Advantages</u>
The other attributes and properties remain unchanged.
<u>Disadvantages</u>
Risk of re-identification:
- possibly on the basis of other properties;
- by combining them with different datasets.</td></tr>
<tr><th colspan="2"><b><i>RECORD SUPPRESSION, OR DELETION OF DATA PERTAINING TO CERTAIN CONSPICUOUS DATA SUBJECTS</i></b></th></tr>
<tr><td><u>Concept</u>
In this process, all data of persons who fall outside certain boundaries, the outliers, are deleted. This is to prevent the individuals who deviate from the mean from being easily identified.
<u>Example</u>
Following a test, the results are made public without name, but with indication of place of residence.
Most participants are from municipalities near where the test was taken, but some participants are from other municipalities. If the scores of participants from other municipalities are published, others who know they participated can easily ascertain what score they achieved.
Consequently, the scores of these participants are not published.</td><td><u>Advantages</u>
The other attributes and properties remain unchanged.
<u>Disadvantages</u>
The data of individuals who deviate from the norm are removed, meaning that the dataset is no longer accurate. For example, the mean and median will be affected.
Risk of re-identification:
- possibly on the basis of other properties;
- by combining them with different datasets.</td></tr>
</table>

<table>
<tr><th colspan="2"><em>CHARACTER MASKING, OR REMOVING CERTAIN CHARACTERS</em></th></tr>
<tr><td>Concept
Certain characters are masked to prevent identification.
Example
Test results are published per post code. Only a few participants come from certain post codes, meaning that they could be identified. However, their post codes only differ from the other postcodes in the last 2 digits.
If only the first 2 digits of the post codes are published, it is still visible which region the participants come from, but no longer from which city/municipality.
The post codes of participants from the Belgian cities Gent (9000) and Gentbrugge (9050) for example are then displayed as "90xx" or as "9xxx". Participants from Leuven (3000) and Kessel-Lo (3010) are displayed as "30xx" or "3xxx".</td><td>Advantages
The other attributes and properties remain unchanged.
Disadvantages
The results are less accurate.
Risk of re-identification:
- possibly on the basis of other properties;
- by combining them with different datasets.</td></tr>
<tr><th colspan="2">GENERALISATION</th></tr>
<tr><td>Concept
The data is generalised, converted to certain categories, meaning that individual data is no longer available.
Example
Test results are published, but instead of age, residence and score, the following are published:
- the age category: "26-30" instead of 27;
- the province: "Flemish Brabant" instead of "Leuven"
- the score group: "76-80%" instead of "77%".</td><td>Advantages
There is no longer any exact data in the dataset and there is less chance of re-identification.
Disadvantages
The results are less accurate.
Risk of re-identification:
- possibly on the basis of other properties, for example in the case of outliers;
- by combining them with different datasets.</td></tr>
<tr><th colspan="2"><em>SWAPPING OR SHUFFLING OF DATA</em></th></tr>
<tr><td>Concept
All personal data remain in the set, but they are swapped so that a set of linked data (a record) no longer relates to the same person.
Example
A dataset contains the name, first name, post code, year of birth and annual income of the data subjects. These are shuffled so that the different data belonging to one person are no longer together.
If the following data are part of the set:
- Janssens, Lenka, 9000, 1987, 30,000 EUR;
- Achmar, Petra, 3000, 1979, 35,000 EUR;</td><td>Advantages
All data is retained.
Disadvantages
Only possible insofar as the intended use allows the different parameters to be shuffled.
Risk of re-identification:
- possibly on the basis of other properties, for example in the case of outliers;
- by combining them with different datasets.</td></tr>
</table>

- Duchateau, Kim, 1000, 1992, 25,000 EUR. Then the data will look like this: - Janssens, Petra, 9000, 1987, 35,000 EUR; - Achmar, Kim, 1000, 1979, 30,000 EUR; - Duchateau, Lenka, 3000, 1992, 25,000 EUR. All data are still present in the dataset, but it is no longer possible to determine which data belong together.	

5. REDUCE THE READABILITY AND USABILITY OF DATASETS FOR THIRD PARTIES

Various **techniques** can be applied to **reduce** the **readability and usability of personal data** by third parties. This makes it more difficult to identify the data subjects and/or certain personal data in a dataset. Consequently, there is less risk of damage to the data subjects, for example in the event of a data breach.

Below we discuss several useful **privacy enhancing technologies** (PETs), to reduce the readability and usability of datasets.

The above-mentioned methods of anonymising personal data are also PETs. When they are applied 'imperfectly' and consequently anonymisation is not actually achieved, they still help to reduce the usability of the data.

DIFFERENTIAL PRIVACY OR PERTURBATION	Applying differential privacy means adding 'noise' to a dataset, which then makes it difficult to determine which personal data is real and which is not.
PSEUDONYMISATION	The GDPR explicitly encourages pseudonymisation as an essential technique for both data minimisation, data protection by design, and data processing security.[92] As discussed earlier, pseudonymisation is considered by some to be a limited form of encryption.[93] Applying pseudonymisation means that the original dataset can continue to exist in the background, but is made more difficult to access. When it is consulted, by hiding certain parameters, only a pseudonymised version can be shown to the user. The pseudonymisation has no effect on the actual dataset, only on its accessibility. Pseudonymise personal data as soon as possible and store the keys in a separate location.
MAKING PERSONAL DATA UNREADABLE	Personal data is made unreadable for people, but remains readable for the computers that use it. This is the case, for example, when personal data are stored in feature vectors.[94]This does not rule out identification and recognition of the personal data, but it does make it more difficult. Depending on the method of application, this technique can be combined with pseudonymisation and encryption, depending also on the prevalence or universality of the encryption used.
ENCRYPTION AND HOMOMORPHIC ENCRYPTION	Personal data is converted into a format that is not readable without the required key. Encryption is also explicitly mentioned in the GDPR on several occasions, including in relation to security of processing[95] and data breaches[96].

[92] See, inter alia, Articles 4(5), 6(3), 25(1), 32(1), 40(2) and 89(1) of the GDPR.

[93] See also section 3.3 on anonymisation and pseudonymisation.

Feature vectors are vectors that contain the characteristics of a parameter.[94] For example, the biometric data needed for face recognition is stored in its mathematical representation. These vectors still allow identification and are personal data, but they are meaningless when seen with the naked eye.

[95] Art. 32(1)(a) GDPR.

[96] Art. 34(4)(a) GDPR.

	Encrypted data are still personal data, given that identification and recognition of the personal data is possible when the data is decrypted. In applying homomorphic encryption, it is possible not only to keep the dataset encrypted, but to train an AI system directly on the encrypted data, without the AI system having access to the non-encrypted data.[97]

6. REDUCE THE VOLUME OF (CENTRALLY) REQUIRED PERSONAL DATA IN THE TRAINING PHASE

There are also privacy enhancing technologies to reduce the volume of (centrally) required datasets in the **training phase**.

DATA MINIMISATION AS A PRODUCT REQUIREMENT	Impose minimal data protection as a product requirement when developing, purchasing or commissioning the development of an AI system that will still need to be trained.
SELECTION OF THE PARAMETERS IN A DATASET	Evaluate which parameters present in a dataset are necessary for the training process and delete the other parameters before using the dataset.
FEDERATED LEARNING	Federated learning makes it possible to train on different databases on local devices, without the personal data leaving the local device. An AI system then trains within different data environments containing personal data. Only the insights learned leave the local environment. The personal data is not shared with the central AI system. This technology is used, for example, in text prediction on smartphones. On every smartphone, the functionality of the AI system is trained. The input on which the system trains does not leave the devices. The insights are however shared, so that the central AI system and its functionality are improved across all devices. This technology is also used in medical research where different patient files need to be used for training. One point to bear in mind for federated learning is that no personal data may be derived from the insights shared and it must not be possible to (re)identify the data subjects. This risk increases the more complex the model is and the more specific the parameters that are used.
GENERATIVE ADVERSARIAL NETWORKS (GANS)	By using a GAN, or Generative Adversarial Network, the required volume of (personal) data for training an AI system is reduced. This involves training two neural networks. One network acts as a generator, the other as a discriminator. The generator will attempt to generate new data, similar to the initial data, based on a dataset. The discriminator will attempt to distinguish the real data, namely those from the original dataset, from the data generated by the generator. The generator will attempt to mislead the discriminator and try to create new data that the discriminator cannot distinguish from the original data. As a result, the generator will improve the quality of the output based on the feedback from the discriminator.

[97] For more information on encryption and the techniques used, see section 4.3. on security of processing.

	The discriminator improves its ability to recognise the non-original data. Although GANs still require fairly large amounts of data in order to be trained, they still allow an AI system to be trained with a more limited initial dataset, supplemented with data obtained via a GAN. Using datasets that are too small runs the risk of insufficient data diversity and therefore inherent reasoning errors or bias slipping into the system.
TRANSFER LEARNING	In transfer learning, an AI system is not trained on (personal) data, but it learns from another AI system that has already been trained. The system itself therefore no longer processes personal data, but somewhere up the chain there will have been some training with personal data. Of course, it is important that the existing AI system from which the learning is done is reliable and does not contain any bias.

7. *REDUCE THE VOLUME OF (CENTRALLY) REQUIRED PERSONAL DATA IN THE USAGE PHASE*

Finally, there are also a number of possibilities to reduce the volume of (centrally) required personal data in the **usage phase** such as:

- data minimisation as a product requirement;
- minimal data protection as a product requirement when purchasing or commissioning the development of an AI system;
- privacy preserving queries (P2Q).

E. An effective and enforced data policy: restrict access to personal data

It is also required to set up an effective data policy in which a **clear and logical definition of roles** ensures, among other things, that:

- as few people as possible have access to the 'raw' personal data;
- there can only be access to personal data with individual accounts and such access is logged;
- for the consultations of the personal data where access to (all) 'raw' personal data is not required, the personal data displayed are limited and/or visually pseudonymised.

APPLICATION

In both e-commerce and recruitment, it can be ensured that, by default, a person who has to call a customer for an appointment only has access to the necessary data in this regard, for example the name, telephone number and reason for the appointment.

4.3. Which measures to consider when securing the processing of personal data by AI systems?

ESSENCE
Personal data can only be properly protected if measures are taken to ensure their integrity and confidentiality. These measures must be technical to protect the infrastructure on which the data are processed and the

data themselves. In addition, organisational measures are required to ensure that individuals within an organisation apply the required measures correctly, handle personal data appropriately and are aware of the importance of data protection.

IN ACTION

- ✓ Establish a data security policy.
- ✓ If not already in place, draw up a record of processing activities and use this as the basis for the considered management of personal data and for the evaluation of data flows.
- ✓ Take the necessary technical measures to protect the ICT infrastructure and personal data against both intentional and accidental or unintentional threats.
- ✓ Implement effective access control and authentication for ICT systems, specific environments containing personal data, and buildings.
- ✓ Implement an enforceable role and authorisation policy that stipulates who gets access to which environment and which personal data.
- ✓ Monitor the ICT environment and access to personal data.
- ✓ Inform and raise awareness among staff within the organisation and establish binding guidelines for the use of the ICT infrastructure and the handling of personal data.
- ✓ Make clear agreements with suppliers and processors about security and compliance with the GDPR.
- ✓ Document all efforts made including the related considerations.

A. Technical and organisational security of the complete environment in which personal data are processed

Personal data can only be properly protected if they are **effectively secured**. There is no point in drawing up rules, guidelines or a policy if anyone who wishes to can gain access to the personal data. 'Security of processing' is one of the cornerstones of the GDPR.

On the one hand, this refers to **technical measures** such as implementing encryption, a firewall or password control. On the other hand, **organisational measures** are also required, such as imposing certain obligations on staff and subcontractors. These measures are intended to prevent the personal data being accidentally or unlawfully:

- shared with or exposed to third parties or persons who should not have access to them (whether in bad faith or not);
- lost or destroyed;
- changed.

Here too, these measures need to take into account the **state of the art**, the **costs of implementation, the context and the risks** for the individuals whose data are being processed. Security is therefore also a dynamic obligation that is **risk- and context-based** and that can also evolve within the same organisation.[98]

The GDPR stipulates that the security measures must include the following, where applicable:[99]

- pseudonymisation of personal data;
- encryption of personal data;
- the ability to ensure the ongoing confidentiality, integrity, availability and resilience of processing systems and services;[100]
- the ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident;

[98] Art. 32(1) and 32(2) GDPR.
[99] Art. 32(1)(a) to (d) GDPR.
[100] Here reference is made to the so-called *CIA triad*, or Confidentiality, Integrity and Availability.

- a process for regularly testing, assessing and evaluating the effectiveness of technical and organisational measures for ensuring the security of the processing.

Adherence to an **approved code of conduct** or having certain processes **certified** can be used as *elements* to demonstrate that the processing is secure. Even then, the adequacy of the measures taken must be assessed in concrete terms.

The controller and the processor are responsible for the security of the processing by any person who processes personal data under their authority or on their instructions.

B. Risk-based approach

The security of processing cannot be viewed in isolation from the risk-based approach of the GDPR, as emphasised in the relevant provisions.[101]

By restricting access, protecting personal data from the outside world and making the data unusable by third parties, the risk that too deep an insight into the privacy of data subjects is possible is also reduced here.

Security of processing is also the most important obligation in **reducing** both the **risk of a data breach and the damage in the event of such a breach**. This applies to data breaches resulting from both unintentional and malicious actions, as well as data breaches caused by both internal and external threats. If a data breach nevertheless occurs, the number of persons affected, the usability of the leaked personal data and thus the extent of the damage to these individuals is significantly reduced.

C. Data security policy

In order to properly organise the security of the processing in an organisation that develops or works with AI systems, it is necessary to set up a **data security policy**. The first step in this regard is to conduct a comprehensive internal analysis of the risks and requirements, and work out how these can be reduced and respectively implemented. As discussed above, these include both organisational and technical risks and measures.

The main **(interrelated) domains** are:

ENCRYPTION AND PSEUDONYMISATION
Data is rendered unreadable if unauthorised persons gain access or if storage media (such as USB storage, hard drives or laptops) fall into unauthorised hands.
CONFIDENTIALITY
It must be defined who should have access to what. Appropriate access control and restrictions are also set up (both physical and electronic). Confidentiality must be guaranteed both with regard to unauthorised internal staff and third parties.
INTEGRITY AS REGARDS ERRORS
The integrity of data can be inadvertently compromised by errors made by employees, systems or external parties. These errors must be detected and prevented.
INTEGRITY AS REGARDS INTENT
Integrity can also be compromised intentionally, when people or systems wish to make the data unavailable, destroy it or change it in bad faith.

[101] Art. 32(1) and 32(2) GDPR.

Resilience and resistance to incidents
In the event of an incident, the proper functioning of the network and infrastructure (such as network, servers, laptops or buildings) can be disrupted. The availability of these elements may be interrupted for staff. As such, measures must be taken to limit the disruption to an organisation's operations in the event of an incident, for example, by providing backups, backup structures, business continuity plans, and incident-response plans.
Incident resolution
Incidents can never be completely ruled out. An organisation must be able to respond quickly to incidents, resolve them and, if necessary, make the required alerts in the event of a data breach.
Evaluation
The adequacy of the security measures must be evaluated and proactively monitored. Learn from mistakes and exercises, and apply continuous improvement.

Setting up a data security policy requires a **multidisciplinary and interdepartmental approach**, as well as management support. On the basis of the **internal data security policy**, an **external data security policy** can also be drawn up: this is one or more documents which outline the (non-confidential) key points of the data security policy. These can be used to make clear to suppliers and partners what is expected in this area, and also to inform customers, the government and the public.

A number of measures that help to safeguard the above-mentioned domains in a cross-cutting manner are explained below.

D. Data management and data mapping

In order to set up an efficient data security policy for personal data, it is necessary to know **what types** and **volumes** of data are processed, **how** the data flows, **what processing** is done with it and **who needs access** to what data. A properly drawn up **record of processing activities**[102] is necessary in this regard.

Based on this information, the optimal data flows can be determined, which will make it possible to organise and monitor access to and protection of these data in the most efficient way, and limit the exposure as much as possible. Of course, the other applicable principles such as data minimisation must also be taken into account in this regard.[103]

E. Technical measures

Technical measures are **necessary** to adequately protect personal data.

Below are several measures that must always be part of a security policy in an organisation that processes personal data using an AI system. In view of the non-technical nature of this contribution, the concepts are indicated without going into further detail regarding the technical properties and possibilities.

For each of the techniques listed below, it is always necessary to take account of the **state of the art** and **current market practice** in order to identify the most appropriate techniques for providing adequate protection.

[102] Art. 30 GDPR.

[103] See also section 4.2 on data minimisation.

STORAGE OF PERSONAL DATA	**Intelligent organisation of storage** Ensure that personal data and sensitive data are stored in a logical and user-friendly manner in the appropriate place and that no unnecessary copies are made (or can be made). **Encrypt personal data** By encrypting personal data, the risk of damage in the event of a data breach is significantly reduced because third parties are, in principle, unable to read them. Given that encryption standards are subject to significant inflation and evolution, it is crucial to take this into account. In any case, it is recommended to encrypt servers, individual computers, and any other storage medium (such as USB sticks and smartphones) and to require that only encrypted devices can be used within the organisation, even if users are allowed to use their own devices ('bring your own device'). **Backup** Make sure regular and usable backups are kept in separate locations and are not directly connected to your own network. Test at regular intervals whether the backups are effective and usable. **Data tagging** Tag personal data to facilitate monitoring of data flows and ensure that the right actions are taken at the right time. **Set up a data retention policy** Set up enforced data retention policies where possible in accordance with data minimisation and storage limitation requirements.[104]
ACCESS TO SYSTEMS AND DATA	**Implement effective access control and authentication** Digital access, both on site and remote, to the ICT environment, personal data and the buildings where services are provided must be subject to individual authentication of the user. This makes it possible to determine who gets access to which data, depending on the role requirements. Access to the actual personal data is limited to persons who have an actual need to access them ('need-to-know'), with individual access. Others will only have access to pseudonymised or anonymised personal data. This prevents unauthorised access to data and limits the risks in the event of both external threat*s (e.g. hacking or phishing) and internal threats (e.g. a disgruntled employee copying company data).* It also allows for faster detection of potentially suspicious behaviour and causes of data breaches. An essential aspect in this regard is that a secure and enforced authentication policy is applied whereby (i) unique, strong and regularly changing passwords are mandatory, (ii) two- or three-factor authentication is applied where useful, (iii) the authentication requirements become stricter depending on the access level of the user and (iv) additional authentication may be required when an individual enters a sensitive environment or logs in remotely.

[104] See also section 4.2. on storage limitation.

	Physical access control should not be overlooked here either: on the one hand, access to the work environment must be closed to external parties, and on the other hand, access to places where local servers are located, for example, must be restricted. In this regard, make sure that access authorisations are immediately deactivated when an employee or service provider leaves the organisation. **Protect systems from the outside world** Protect the network environment from external access by third-parties and unauthorised persons, with detection and the possibility to automatically take protective measures. **Destroy physical (paper) and digital storage media**
MONITOR THE NETWORK ENVIRONMENT AND BUILDINGS	**Actively monitor the ICT environment and the buildings for**: - (attempted) security breaches, both via the network and via other channels (e.g. e-mail); - suspicious access, such as a logon from another country; - suspicious behaviour, such as login attempts at an unusual time or repeated failed authentication attempts; - large movements of data; - malicious applications such as viruses and malware; - data breaches. There must be specific consideration for this, so that the right people are alerted and if necessary the notification obligation can be fulfilled quickly. Depending on the risk linked to certain incidents, consequences can then be linked automatically, such as (i) requiring (once again) the user to log in or to provide the second authentication factor, (ii) displaying a warning (e.g. in case of downloading personal data), (iii) alerting certain services within the organisation, (iv) restricting the access of the user in question or (v) shutting down certain processes.
COMMUNICATION	**Security** Secure emailing and other communications with encryption, spam and phishing protection. Allow personal data to be shared only via your own platform and via temporary and/or personalised hyperlinks. **Connection** Also make sure that communication with and logging in to the systems is always via a secure connection.
SOFTWARE POLICY	Implement an enforced software policy so that only trusted and approved software is used on systems with access to the network. Make sure software is always up-to-date.

F. Awareness and training of all persons with access to personal data

In order to secure the processing of personal data, it is not enough to take technical measures. These measures must also be **understood and supported** by the persons who apply them. Indeed, there is no point in putting five locks on a door if the door is always left open. Whenever reference is made below to staff members, this should be understood as meaning everyone within the organisation who processes personal data and/or has access to the network environment, i.e. also external consultants and administrators within the organisation.

The following points are particularly important in this regard.

ASSIGNED ROLES AND ACCESS AUTHORISATIONS	Ensure that access to systems, processes and personal data is appropriate to the tasks that the individuals concerned have to perform and that they understand why these rights and restrictions are applied in that way.
POLICY	Establish rules for the proper use and handling of the ICT infrastructure and personal data. Ensure these rules meet the labour law requirements to be enforceable. The following guidelines are necessary in this regard (whether or not they are part of a single document or multiple documents): - ICT policy: this stipulates the safe approach to ICT equipment, passwords, internet use, e-mail, etc. - data policy: this specifically defines how personal data and sensitive data must be handled. In a context where AI systems are used, there must be a specific focus on the use of personal data in combination with such systems. - BYOD ('Bring Your Own Device') policy: this sets out the conditions under which personal devices may be used within a network environment or, generally, for work purposes.
TRAINING	Organise general and specific training for staff on data security, data protection, ICT security and use of personal data in an AI context. The intensity of the training and the number of sessions depend on the extent to which the staff members in question process data, have an influence on data processing through their tasks and/or are responsible for ICT security.
AWARENESS-RAISING	Rules are only correctly applied if they are supported by the people who apply them. Therefore, make sure that the staff members are also made aware of the importance of data protection and are aware of the risks. This is achieved, among other things, by: - getting the board and management to set a good example and emphasise the importance the organisation attaches to data security; - organising both fun and serious information campaigns, either in general or by topic/theme; - setting up test campaigns, for example by sending out in-house phishing mails.

G. Agreements with suppliers and processors

Make **clear agreements** with suppliers of goods, services and software that can affect the security of the infrastructure and the processing of data. Oblige them to comply with the requirements of the GDPR and impose minimum security requirements. Make sure they are liable for any damages they cause, if they are not in compliance.

If suppliers process data as a processor, a **data processing agreement** must always be entered into.[105]

[105] See also section 3.5on the roles and responsibilities of the controller and processor in an AI context.

H. Document

Under the general 'accountability' requirement of the GDPR, it must be possible to prove compliance with the security of processing obligation.[106] It must therefore be possible to prove that the required **measures have been taken and that they are effective**.

In this regard, it is necessary to **document**, among other things:

- what technical measures have been taken and why;
- how past problems and incidents were taken into consideration;
- what training has been given to staff and why they needed to follow it;
- what awareness-raising campaigns were implemented; and
- that the necessary agreements with all suppliers and processors have been made and implemented.[107]

4.4. When does a DPIA need to be conducted for the processing of personal data by AI systems?

ESSENCE

A DPIA is a tool to identify in advance the data protection risks of a data processing operation and to be able to subsequently take measures to mitigate the identified risks.

AI systems often involve both new technologies and complex and unexpected outcomes relating to personal data, which may require a DPIA. Taking into account the direction of the White Paper on AI[108], it seems advisable that companies developing AI systems for application in the context of recruitment or e-commerce should already consider conducting DPIAs.

Performing a DPIA also allows organisations to demonstrate that the processing of personal data by an AI system is proportionate.

IN ACTION

- ✓ Think about what form of DPIA the organisation wishes to implement and create a DPIA template (see below for a blueprint of step-by-step plan).
- ✓ Examine which processing operations are considered high risk.
- ✓ Investigate whether the organisation is a processor or a controller with regard to the intended processing.
- ✓ Always perform the DPIA as early as possible in the life cycle of an AI system.
- ✓ Examine whether the processing of personal data is necessary and consider whether its use is proportionate to the ultimate purpose.

A. General

Every organisation that processes personal data must assess whether there are any risks involved. If an organisation suspects that an AI system is likely to pose a **high risk** to the **rights and freedoms of natural persons**, it must conduct a DPIA.[109]

[106] Art. 5(2) GDPR. The accountability principle is one of the basic principles of data processing. It requires controllers and processors to be able to demonstrate that they have taken steps to comply with the obligations under the GDPR.

[107] European Data Protection Board, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, 13 November 2019, p. 6, no. 15.

[108] The White Paper on AI states that: "In light of its significance for individuals and of the EU acquis addressing employment equality, the use of AI applications for recruitment processes as well as in situations impacting workers' rights would always be considered "high-risk" and therefore the below requirements would at all times apply. Further specific applications affecting consumer rights could be considered".

[109] Art. 35(1) GDPR.

A DPIA is intended to **describe the processing of personal data, assess** its **necessity and proportionality**, and help manage the **associated risks** to the rights and freedoms of natural persons by assessing these risks and defining measures to address them.

Although a DPIA is only mandatory for processing with a **likely high risk**, it is recommended to perform one in **other situations as well**. Indeed, it is a useful tool that helps organisations to comply with data protection legislation.

There are **various methods** for performing a DPIA. There are no specific formal requirements , but the DPIA must include as a minimum:

- which data will be processed, in what way they will be processed and why they will be processed;
- why these processing operations are necessary and their proportionality;
- what measures are taken to address the risks (such as safeguards and security measures).

If the planned processing operations were to present a high risk and the organisation is unable to provide measures to mitigate this risk, these processing operations must be **submitted to the competent DPA in advance**. This is possible via this form.[110]

To assess the level of risk, a DPIA must consider both the **likelihood** and **severity of any impact** on individuals.

B. Determining the risk

The GDPR stipulates that a DPIA is required in the following **non-exhaustive cases**:

- systematic and extensive profiling with significant consequences;[111]
- processing on a large scale of special categories of personal data;[112] or
- systematic monitoring of a publicly accessible area on a large scale[113].

To **determine whether a processing operation is high risk** and therefore may require a DPIA, the following criteria must be considered:

1. Evaluation of personal aspects relating to natural persons

This includes profiling and the making of predictions, in particular about characteristics relating to the professional performance, economic situation, health, personal preferences or interests, reliability or behaviour, location or travel of the data subject.

One example is a company that collects information about the visitors to a web shop (e.g. purchase history, browsing behaviour and other online information) and uses this information to create profiles of these individuals (or has them created by an AI application) so that they can automatically offer personalised advertisements.

2. Automated decision-making

This involves making decisions by technological means without any human intervention. This may also include profiling.[114] If such a decision results in someone facing legal consequences or they are otherwise significantly affected, a DPIA is mandatory.

In any event, recruitment based purely on an algorithm, i.e. without human intervention, will have a significant impact on a person, partly because it decides whether or not a person gets a job and because there may be discrimination.

[110] See also: https://www.dataprotectionauthority.be/publications/act-of-30-july-2018.pdf

[111] See also section 5.4 on profiling.

[112]See also section 3.2 on special categories of personal data

[113] Art. 35(3) GDPR.

[114] See also section 5.4 on profiling.

3. 'Sensitive' data

These are special categories of personal data[115] such as information about political opinions, criminal offences, sexual orientation or medical data. It can also be data that are generally considered to be privacy sensitive such as electronic communication data, location data and financial data.

Through a recruitment process, there is a risk that sensitive data such as an individual's financial situation, union membership or medical data may be collected either directly (through a job interview) or indirectly (through social media).

4. Large-scale data processing

The GDPR does not define large-scale data processing. The WP29 already recommends taking into account the following factors:

- The number of data subjects concerned;
- The volume of data and/or the range of different data items being processed
- The duration, of the data processing activity
- The geographical extent of the processing activity.

5. Combining data from different sources

This involves matching or merging data sets. For example, data collections resulting from two or more processing operations for different purposes and/or carried out by different controllers, in a way that does not correspond to the reasonable expectations of the data subject.

For example, an online retailer wants to supplement the current customer base with essential information, so that it can respond to the needs of customers more effectively and offer them personalised ads. To this end, it solicits a company that specialises in this area. By merging its customer base with that of the third party, it has an 'enriched' customer data base.

6. Vulnerable data subjects

A DPIA may be necessary because there is an unequal balance of power between the data subjects and the controller. As a result, these data subjects are not able to freely give their consent to, object to or exercise their rights with regard to the processing of their data. Vulnerable data subjects include children, workers, patients and the elderly.

7. Use of new technologies

The GDPR clearly states that a DPIA may be required when new technology is applied. Indeed, the use of new technology may involve new forms of data collection with potentially high data protection risks (cf. COVID contact and health apps).

Indeed, the personal and social consequences of using a new application of AI technology may still be unknown. A DPIA can help to assess and address the possible risks and, for example, determine what additional information needs to be given to the data subjects so that they can assess the impact of the processing of their data. As an AI system acts more autonomously, has more leeway to make decisions and is even able to select data sources on its own, it becomes more important to thoroughly analyse the possible consequences of this autonomy.

8. Exclusion

This is about data processing that has the following effects on individuals:

- they cannot exercise a right;
- they are unable to use a service; or

[115] See also section 3.2 on special categories of personal data

- they cannot conclude a contract.

One example of this is an employer screening applicants based on a database of CVs and references before deciding whether to hire them or not.

In most cases, an organisation can assume that a DPIA needs to be performed for a processing operation that meets two of the above criteria. In general, the WP29 assumes that the greater the number of criteria that a processing operation meets, the more likely it is to present a high risk to the rights and freedoms of data subjects. Consequently, a DPIA is required irrespective of the measures that the controller may (still) take.

C. List of the Belgian DPA

According to the GDPR, each national supervisory authority needs to establish and publish a list of the **kind of processing operations which are subject to a mandatory DPIA**.[116]

The Belgian DPA has published a **non-exhaustive list** of the following processing operations for which a DPIA is required[117] :

1.	When the processing uses **biometric data for the unique identification of data subjects** who are in a public area or in private areas accessible to the public. AI systems that use facial recognition, for example, may fall under this.
2.	When personal data is collected from third parties to be taken into account in the decision to **refuse or terminate a specific service agreement with a natural person**. One example is financial institutions that use algorithms to search for information about a customer's creditworthiness.
3.	When **health data** of a data subject are **automatically** collected using an active **implantable medical device.** This includes, for example, any active ('smart') medical device that is designed to be wholly or partially implanted in the human body.
4.	When data are **collected from third parties on a large scale** in order to analyse or predict the economic situation, health, personal preferences or interests, reliability or behaviour, location or movements of natural persons. For example, a web shop that uses an AI system to collect as much information as possible about a person across different channels, in order to make as detailed predictions as possible.
5.	When **special categories of personal data or data of a highly personal nature** (such as for example data about poverty, unemployment, involvement of youth welfare or social work, data about household and private activities or location data) are systematically exchanged between multiple processors.
6.	When there is a situation of **large-scale processing of data** generated by **devices with sensors that send data over the Internet or through another medium** (Internet of Things or IoT applications such as smart televisions, energy meters and household appliances or smart cities) and this **processing serves** to analyse or predict the economic situation, health, personal preferences or interests, reliability or behaviour, location or movements of natural persons.
7.	When there is a situation of **large-scale and/or systematic processing** of telephony, internet or other communication data, metadata or location data of, or traceable to, natural persons (for example WiFi tracking or processing of location data of public transport passengers) when the processing is not strictly necessary for a service requested by the data subject.
8.	When there is a situation of **large-scale processing of personal data** whereby the behaviour of natural persons is systematically observed, collected, recorded or influenced via automated processing, including for advertising purposes. For example, profiling based on existing customer data to serve personalised ads.

[116] Article 35(4) GDPR.

[117] See: (Dutch or French only) Decision of the General Secretariat No. 1/2019 of 16 January 2019, Belgian Official Gazette March 22, 2019, 28512-28514 (www.gegevensbeschermingsautoriteit.be/sites/privacycommission/files/documents/01_2019_AS.pdf).

It **may also be necessary to perform a DPIA for existing processing operations**. This is in particular the case when:

- the risk of processing changes, for example, because a new technology is used, the nature of the previously identified processing risk requires regular reassessment, or personal data is used for a different purpose;
- the organisational or social context changes, for instance whencertain automated decisions have become more important or because new categories of data subjects become vulnerable to discrimination. Take recruitment, for example, where manual profile searching to perform job matching is replaced by an AI system that does this task automatically without human control.

APPLICATION

When a web shop collects personal data from customers in order to arrange deliveries, this is a case of 'existing processing' which does not require a DPIA. New processing of the same data for which a DPIA is required would occur if the organisation starts using this customer database for new purposes, such as profiling.[118]

However, a **DPIA is not required** if:

- the processing is unlikely to present a high risk;
- the processing was already verified by the DPA and the processing has not changed in the meantime;
- the risks arising from the processing have not changed.

D. Time line for carrying out the DPIA

A DPIA must be carried out **prior to the processing** of personal data.[119] In practice, this means that it is best to carry out the DPIA at the earliest possible stage, such as the design phase of the product or process in which data will be processed, even if certain processing is not yet known. By carrying out the DPIA as early as possible, it will also make it easier to meet data protection by design/by default requirements.[120]

It has already been mentioned that performing a DPIA is not a one-off exercise, but an **ongoing process**. The processing operations must be monitored and the DPIA may therefore need to be subsequently updated.

E. Who needs to carry out a DPIA?

The **controller** must carry out the DPIA.[121] The DPIA may be carried out within or outside the organisation, but the controller has ultimate responsibility.

If a **Data Protection Officer** (DPO) has been appointed, the controller must also seek his or her advice.[122] The DPIA report must contain the opinion and the actions taken by the controller. The DPO also supervises the carrying out of the DPIA.[123]

F. Step-by-step plan

Based on the above components, a **concrete step-by-step plan** can be drawn up:

[118] See also section 5.3 on profiling.
[119] Art. 35(1) GDPR.
[120] See also section 4.1 on data protection by design.
[121] Art. 35(1) GDPR.
[122] Art. 35(2) GDPR.
[123] Art. 39(1)(c) GDPR.

1.	Examine and specify as early as possible in the development of an AI system whether a DPIA needs to be carried out, whereby the following questions, among others, are relevant: - What are the objectives of the AI system? - Is a new technology used which could potentially have a major impact on data protection? - Will the AI system lead to certain decisions or actions against individuals in a way that could significantly affect them? The implementation of an AI system to process personal data must be guided by the fact that the system can achieve a specific and legitimate purpose and therefore not merely because the technology is available. By describing the need for using an AI system in a DPIA, an organisation can demonstrate that this purpose could not be achieved in any other reasonable way.
2.	Systematically describe the intended processing operations, indicating in each case the possible legal grounds and with specific consideration for: - Identifying the data flows and stages in which an AI application and automated decisions could impact individuals; - Using as much anonymous data as possible or mapping information flows as the project progresses, if the processing purposes are not yet certain (e.g. due to correlations that are not yet known); - Specifying and recording the roles and obligations of controllers and processors. Where processing by AI systems is entirely or partially outsourced to a third party, all organisations involved must assess whether a joint control structure is in place. If this is the case, they must cooperate with the DPIA.[124]
3.	Map the data protection and related risks, with consideration for: - assessing the necessity and proportionality of the processing in relation to the purpose; - assessing the risks to the rights and freedoms of data subjects; - the fact that drawing up a DPIA is not a one-off action, but rather a 'living' and continuous process that develops during (the evolution of) a project. The following questions may arise in this regard: o Have data subjects been made aware of the use of their personal data? o Can the dataset contain sensitive data? o What are the retention periods of the processed data? o Are the data stored on multiple systems? o Do the systems have appropriate security systems? o Can anonymised data be re-identified? - Methods other than big data analysis for this project. The fact that organisations need to take other legislation into account. For example, the use of an AI system may result in discrimination based on historical data patterns, which may violate anti-discrimination regulations.
4.	Identify and evaluate possible technical and organisational solutions or measures to rule out/mitigate data protection breaches.
5.	Conclude the DPIA.
6.	Integrate the results of the DPIA into the project plan of the AI system. It is important that the people who will be implementing the AI project understand the solutions and measures, why they are necessary and how they can be implemented.

[124] See also section 3.5 about the roles and responsibilities.

This is the final step in the process, but as already mentioned, not the end point. Regular evaluations should ensure that the proposed solutions are working as intended. Moreover, the objectives and applications of the project may change over time. Regular assessments can help determine these changes and check whether the DPIA needs to be adapted.

4.5. What are the obligations when personal data are processed for scientific research or statistical purposes?

ESSENCE

Any processing of personal data for scientific or statistical purposes must comply with the requirements of the GDPR. However, both the GDPR and Belgian legislation[125] provide some exceptions for researchers with regard to, for example, the rights of data subjects, retention periods or the compatibility of further processing.

In Belgium, furthermore, a 'cascade' system applies. In principle, research must be done with anonymised data. Only if this is not possible can pseudonymised or identifying personal data, respectively, be used.[126]

Belgian law also addresses five other topics:

- ✓ Data collection directly from the data subject: the data protection declaration must contain additional information;
- ✓ Indirect data collection / further processing of data: an agreement must be concluded with or a notification made to the original controller;
- ✓ When and how to anonymise or pseudonymise the data: depending on the situation, the original controller or a trusted third party adviser, who is under a confidentiality obligation, must anonymise or pseudonymise the transmitted personal data at some point;
- ✓ Dissemination of data: identifying personal data may not be disseminated unless certain exceptions apply. Pseudonymised data may be disseminated, unless certain legal restrictions apply, or if it is sensitive data;
- ✓ Communication of data: identifying personal data may be communicated/transmitted to an identified third party for scientific or statistical purposes. However, in certain cases they may not be reproducible, unless exceptions apply (see below).

IN ACTION

- ✓ Consider whether the necessary technical and organisational measures are in place to comply with the principle of data minimisation.
- ✓ Consider to what extent the (possible) exercise of data subjects rights under the GDPR threatens to render the achievement of the specific scientific or statistical research purposes impossible, or seriously impair them, and to what extent not or only limitedly having to respond to such requests is necessary to achieve the purposes.
- ✓ Map out in which cases personal data are further processed for a scientific or statistical research purpose within the meaning of the GDPR (e.g. training of AI system for commercialisation). In such cases, inform the data subjects before proceeding with further processing (unless one of the exceptions applies).
- ✓ Identify where it would be useful to keep personal data for longer than the strictly necessary period, insofar as they are only (further) processed for scientific research or statistical purposes. In such cases, appropriate technical and organisational measures must be taken, inter alia, to limit access to and use of such data.
- ✓ Perform a DPIA if required by the GDPR.
- ✓ Include the additional information specifically required by law in the record of processing activities and the relevant data protection declaration.

[125] Act of 30 July 2018 on the protection of natural persons with regard to the processing of personal data.

[126] See also section 3.3 on anonymisation and pseudonymisation in this regard.

- ✓ Enter into an agreement with the third party provider of the personal data that meets the legal requirements if personal data are processed that were not collected directly from the data subject. If no agreement needs to be entered into, send a notification in accordance with the legal requirements to the original controller.
- ✓ Determine under which of the legally described situations the processing falls, to ascertain who needs to anonymise or pseudonymise personal data and when.
- ✓ Do not make non-pseudonymised data public unless it is certain that one of the exceptions applies.
- ✓ Ensure that non-pseudonymised data is communicated in a non-reproducible manner to an identified third party if there is a situation as described in the law.

A. Scientific or statistical purposes

This section discusses whether and when AI developers and researchers can rely on the (national) **exception regime** to the GDPR for processing carried out by them in the context of (scientific) research or for statistical purposes. It is therefore important in the first instance to understand what the GDPR understands under **scientific research** and **statistical purposes**.

SCIENTIFIC RESEARCH

Scientific research is interpreted in a broad manner, and includes for example technological development and demonstration, fundamental research, applied research and privately funded research. It also includes studies conducted in the public interest in the area of public health.

It can be assumed that AI developers, researchers and even users will, in certain circumstances, engage in what the GDPR defines as 'scientific research'**.** For example, training an AI system during the development phase may fall under the technological development of the system. Fundamental AI research, whether privately or publicly funded, may also fall under this exception regime. Training an AI system during the use phase, on the other hand, will not be considered a technological development.

PROCESSING FOR STATISTICAL PURPOSES

Statistical purposes mean any operation of collection and the processing of personal data necessary for statistical surveys or for the production of statistical results.[127] These statistical results may be used for different purposes, including a scientific research purpose.

A statistical purpose means that the result of the processing for statistical purposes is not personal data, but aggregated data. This result and the related personal data must not give rise to any measures or decisions regarding a specific natural person (conversely, a group of persons is allowed).

Again, it can be assumed that AI developers, researchers and even users will, in certain circumstances, carry out processing operations that fall within what the GDPR refers to as 'statistical purposes'. For example, usage (or user) statistics or accuracy analyses may be conducted both during training and during the implementation of an AI system (to the extent that personal data is used in this regard) for the purpose of obtaining statistical results. In addition, the GDPR does not rule out the possibility that such results may subsequently be used for commercial purposes.

[127] Recital 162 GDPR.

B. Applicable GDPR principles

The GDPR contains a number of **important rules** regarding the (initial and further) processing of personal data for scientific research or statistical purposes. The table below provides a summary of these.[128]

1.	Such processing should (in principle) respect the rights and obligations that a data subject enjoys under the GDPR. More specifically, the principle of data minimisation must be guaranteed by technical and organisational measures (e.g. pseudonymisation), to the extent that such purposes can still be fulfilled.[129]
2.	When personal data are processed for these purposes, national law may provide for derogations from the rights of access, rectification, restriction of processing, objection and erasure, in so far as the right referred to is likely to render impossible or seriously impair the achievement of the objectives of that processing. However, this necessity requirement does not apply to requests for erasure.[130]
3.	Where processing for scientific research or statistical purposes simultaneously serves another purpose, the possible exemptions shall apply only to the processing for scientific or statistical purposes.
4.	The GDPR also provides some additional derogations or exceptions: - In principle, further processing for scientific research or statistical purposes should not be considered to be incompatible with the initial purposes (of the initial processing).[131] In such cases, the controller must, in principle, inform the data subject before proceeding with the processing. However, the controller may be exempted from this information requirement (in the context of their research activities) if they process data that they have not collected directly from data subjects. However, this is only the case if the provision of such information proves impossible or would involve a disproportionate effort.[132] Appropriate measures must then be taken, such as making the relevant information publicly available (e.g. on a website, if it sufficiently reaches the data subjects); - Personal data may also be stored for longer periods than strictly necessary where they are processed solely for scientific research or statistical purposes, provided that appropriate technical and organisational measures are taken to protect the rights and freedoms of the data subject (for example, pseudonymisation or restricted access to the data).

For more information on the distinction between **pseudonymised** and **anonymous personal data**, please refer to section 3.3. It is also emphasised that processing carried out for scientific or statistical purposes may give rise to the need to complete a **DPIA**. This is explained in more detail in section 4.4.

C. Applicable Belgian legislation

These provisions of the GDPR are **further supplemented** in articles 186 et seq. of the Belgian Act on the protection of natural persons with regard to the processing of personal data. This concerns both **general and specific additions**.

GENERAL ADDITIONS
The law reiterates that exceptions to certain rights of data subjects are possible in so far as they are likely to render impossible or seriously impair the achievement of the objectives of that processing, and derogations are necessary to achieve those purposes.

[128] Art. 89 GDPR. This article also covers archiving in the public interest and historical research. These categories are not further discussed here.

[129] See section 3.3 on data minimisation in this regard.

[130] Art. 17(3)(d) GDPR.

[131] Art. 5(1)(b) GDPR. This was also confirmed by the Spanish Data Protection Authority.

[132] Art. 14(5)(b) GDPR.

The Act subsequently emphasises the fact that a controller processing personal data for statistical or scientific purposes must still designate a DPO if the processing of the personal data may involve a high risk as referred to in Article 35 of the GDPR.

Finally, the law also stipulates that the controller must add the following elements to its record of processing activities prior to the collection of data for scientific or statistical purposes:[133]

- the justification for the use of data, pseudonymised or otherwise;
- the reasons why the exercise of the rights of the data subject risks rendering the achievement of the purposes impossible or risks seriously impairing them;
- where appropriate, the DPIA if the controller processes sensitive data for scientific or statistical purposes.

SPECIFIC ADDITIONS

It is also important to note that Belgian law introduces three new concepts:

- trusted third party advisers subject to a confidentiality obligation: the natural person or legal entity, de facto association or public authority which pseudonymises the data and which is not the controller for the purposes of processing for archiving or research or statistical purposes;
- communication of data: communication of data to an identified third party;
- dissemination of data: disclosure of data without identification of the third party.

Moreover, the law proposes the following '**cascade principle**' for the anonymisation or not of personal data.

PRINCIPLE	The controller for processing for scientific research or statistical purposes uses anonymous data.
EXCEPTION	If it is not possible to achieve the research or statistical objective with anonymous data, the controller will use pseudonymised data.
EXCEPTION	If it is not possible to achieve the research or statistical objective with pseudonymised data, the controller will use non-pseudonymised data.

The Belgian law specifically addresses **five topics** relating to processing for scientific or statistical purposes.

1. COLLECTION OF DATA DIRECTLY FROM THE DATA SUBJECT

In addition to the information that a controller must communicate in accordance with Article 13 of the GDPR[134], the law requires that the data subject is also informed of:

- whether the data are anonymised or not;
- the reasons why the exercise of the rights of the data subject risks rendering the achievement of the purposes impossible or risks seriously impairing them;

This information must be provided before personal data are collected from the data subject.

As such, if in the applications of e-commerce and recruitment, personal data are collected directly from data subjects and are also used at a later stage for scientific or statistical research purposes, the data protection declaration of the controller must state this information.

2. INDIRECT DATA COLLECTION / FURTHER PROCESSING OF DATA

If personal data are not collected directly from the data subject, the controller must enter into an agreement with the controller for the original processing, i.e. the source of the personal data.

[133] In principle, all controllers and processors must maintain a record of processing activities under their responsibility. Article 30 of the GDPR stipulates which data this record must contain.

[134] See also section 5.1

However, such an agreement does not have to be concluded if:

- the processing relates to personal data that have been made public (for example, through social media); or
- when the law of the European Union, a national law, a decree or an ordinance:
    - gives the controller a mandate to process personal data for scientific or statistical purposes; *and*
    - prohibits the re-use of the data collected for other purposes.

In the event of an exemption for concluding a contract, the controller (the researcher) must notify the intended data collection and processing to the controller of the original processing (e.g., the social media platform).

The above-mentioned agreement or notification must contain the following elements as a minimum:

- in the event of a contract, the contact details of the controller for the initial processing and of the controller of the further processing;
- the reasons why the exercise of the rights of the data subject risks rendering the achievement of the purposes impossible or seriously impairing them;

This agreement or notification must also be attached to the record of processing activities by the controller.

It is possible that companies working within recruitment or e-commerce will rely on external sources (e.g. data brokers) to gather data to train their AI systems with or perform accuracy analyses. As such, they need to be aware that they will have to conclude an additional agreement in order to comply with Belgian law.

3. WHEN SHOULD DATA PROCESSED FOR SCIENTIFIC OR STATISTICAL PURPOSES BE ANONYMISED OR PSEUDONYMISED AND BY WHOM?[135]

Situation 1: In accordance with the 'cascade principle' described above, the controller must proceed to anonymise or pseudonymise the data after they are collected, if the data are collected directly from the data subject.

Situation 2: If the controller already has personal data in its possession (in connection with earlier processing) and wishes to process them itself for scientific or statistical purposes, the controller will anonymise or pseudonymise the data *prior to the further processing*.

This controller may only de-anonymise or de-pseudonymise this personal data if this is necessary for research or statistical purposes, and if necessary following the advice of the DPO, which must be documented.

Situation 3: If a controller transfers the personal data to another controller, the original controller will pseudonymise or anonymise the data prior to communicating it to the controller for further processing.

The controller for the further processing cannot have access to the keys for the pseudonymisation.

Situation 4: If multiple original processing operations are linked, the original controllers shall have the data anonymised or pseudonymised by one of the controllers of the original processing or by a trusted third party adviser bound by a confidentiality obligation prior to communicating the data to the controller.[136]

If one of the original controllers transmits sensitive data in such a situation, only that controller can anonymise or pseudonymise the data (or a trusted third party adviser, subject to a confidentiality obligation) prior to communicating the data to the controller for further processing.

[135] If the relevant controllers have appointed a DPO, the DPO should advise on the use of different pseudonymisation and anonymisation methods, in particular on the effectiveness of the data protection.

[136] The 'confidential' third party adviser must (1) be subject to professional secrecy within the meaning of article 458 of the Penal Code; and (2) be dependent on neither the initial controller nor the controller for further processing.

Only the controller for the initial processing who pseudonymised the data or the trusted third party adviser shall have access to the pseudonymisation keys.

4. (PUBLIC) DISSEMINATION OF DATA

Unless certain legislation imposes stricter conditions on the dissemination of data[137] processed for scientific or statistical purposes, the controller cannot disseminate non-pseudonymised data unless:

- the data subject has given his/her consent; or
- the data has been made public by the data subject; or
- the data is closely linked to the public or historical nature of the data subject; or
- the data is closely linked to the public or historical nature of the events in which the data subject was involved.

Pseudonymised data may be disseminated by the controller, unless this is prevented by certain legislation or if it is sensitive data. However, anonymised data may be disseminated.

5. COMMUNICATION OF DATA

Unless legislation imposes stricter conditions, a controller who communicates/transmits non-pseudonymised data to an identified third party for scientific or statistical purposes shall ensure that the identified third party cannot reproduce the communicated data, except in a handwritten manner, if:

- it is sensitive personal data;[138] or
- the agreement between the controllers of the initial processing and of the further processing prohibits this; or
- such reproduction might compromise the safety of the data subject.

This obligation does not apply if:

- the data subject has given his/her consent; or
- the data has been made public by the data subject; or
- the data is closely linked to the public or historical nature of the data subject; or
- the data is closely linked to the public or historical nature of the events in which the data subject was involved.

If a company involved in recruitment wishes to share obtained personal data (e.g. from a CV) with a third party for scientific or statistical purposes, it must ensure that if sensitive data is also shared, these cannot be reproduced. Conversely, if a company involved in e-commerce receives non-pseudonymised data for scientific purposes, it must ensure that it cannot reproduce any sensitive data.

[137] See above for the definition of dissemination of data.

[138] See also section 3.2 on data minimisation.

5. How can data protection be ensured during the deployment of AI systems?

The following sections discuss a number of issues from the GDPR that are important when using AI systems: transparency obligations, storage limitation, the rights of data subjects and automated individual decision-making, including profiling.

5.1. What transparency obligations does the GDPR impose and what are the specific things to bear in mind in an AI context?

Essence

Under the GDPR, controllers must provide certain information to data subjects. There are both general transparency obligations that apply to all types of processing, and specific transparency obligations that must be complied with for certain processing operations that use AI systems. This applies all the more if automated individual decision-making is involved.[139]

In action

✓ Ensure that the organisation has a privacy statement that contains all the information required by the GDPR and is communicated to the data subjects at the appropriate time;

✓ Consider working with a layered privacy statement, especially if useful information relating to the underlying logic of the AI system needs to be provided;

✓ Consider using visual and interactive techniques to communicate this information to the data subjects in a clear and understandable way;

✓ Identify which processing operations using AI systems involve automated decision making and whether these processing operations entail legal or other significant consequences with respect to the data subjects;

✓ When developing AI systems, try to use an 'explainability by design' approach and strive for the most transparent design of AI systems possible;

✓ Inform data subjects as soon as they interact with an AI system that involves automated decision-making;

✓ Think about what information the organisation wants to communicate if useful information is to be provided about the underlying logic of an AI system, and how this information will be provided;

✓ Inform the data subject about the intended or expected processing by the AI system and the consequences that the automated decision making may cause to the data subject, using tangible examples.

The private life of individuals is becoming increasingly transparent to the organisations that use AI systems (and more broadly big data analytics). Conversely, these systems are often themselves characterised by a lack of transparency, or even opacity, vis-à-vis the data subjects and supervisory authorities. Nevertheless, the GDPR lays down transparency obligations that must also be complied with in an AI context.

A. External and internal transparency

In view of the discussion of the transparency obligations under the GDPR, we will first briefly consider the distinction between **external** and **internal transparency**. This is a distinction that does not in itself explicitly follow from the GDPR, but in practice can clarify several elements.

External transparency

This term refers to the transparency that a controller must provide to the outside world as regards the personal data that it processes. In other words, it comes down to translating what is happening on a technical level in an AI system into terms and reasoning that are understandable to the data subject. Such transparency

[139] See also section 5.4 on automated individual decision-making.

is important not only for the data subjects themselves, but also for other stakeholders such as the media, interest groups in the social field and supervisory bodies.

In articles 12-14 the GDPR specifies what information must be provided. This is therefore not a detailed explanation of what an AI system can do and what it does actually do, but rather and primarily about managing the expectations of third parties.

INTERNAL TRANSPARENCY[140]

This refers to the transparency, relating to the operation of an AI system, that should ideally be ensured within an organisation. Not only do IT teams need to be able to understand what the AI systems being used are capable of and how they work, but also product and account managers, the data protection officer, in-house legal counsels or executives need to have relevant information so that they use these systems in an appropriate and informed manner.

The GDPR refers to the term 'accountability'. In this context, it is recommended to document in detail how an AI system, for example, processes data, which technical principles it is based on, (a description of) what data is processed, who played what role at what point in the development process, and which training methodologies were applied. This information can be included in internal guidelines, regulations or memos addressed to the different functions within an organisation, and provide information on the operation of the AI systems used.

The following parts will explore the first section and examine what information must be incorporated in a privacy statement according to the GDPR.

B. Transparency obligations under the GDPR

The first section discusses the **general information obligations** that must be complied with under the GDPR. The next part considers in more detail the **information requirements specifically in relation to AI systems** that could be clarified. It is important to realise in an AI context that the transparency obligations cover all **the phases of the data processing**, i.e. the training and testing phase as well as the application phase of the AI systems.

In an AI context, transparency must allow the data subjects to understand the implications of such AI systems. Transparency is aimed at both data subjects and controllers. More specifically, transparency relates to accurate information about the **actual possibilities and limitations** of AI systems so that false **expectations** among data subjects and incorrect interpretations of results are avoided. Transparency also includes providing information about the context of the processing, the involvement of third parties, etc.

The information obligations set out below provide a non-exhaustive list of the general transparency obligations. Several documents have already been published that explain this in more detail.[141]

General information obligation

Articles 12, 13 and 14 of the GDPR contain the main general transparency obligations that controllers must comply with. They are briefly discussed below.

ARTICLE 12 GDPR.

[140] This section, and the related explainability question, will not be discussed in detail due to its more technical nature and the fact that it is not explicitly imposed by the GDPR. For more information, see the 'Explain AI' project of the British ICO and the Alan Turing Institute, or the report 'Robustness and Explainability of Artificial Intelligence' by the European Commission.

[141] For more information, see for example the rather detailed and authoritative guidelines on transparency of the former Article 29 Working Party (WP29).

Article 12 GDPR lays down in a general manner that the data subject must receive the information in a concise, transparent, intelligible and easily accessible form, using clear and plain language.

This means that information must be presented in a way that allows an average educated person to easily find and understand the information, whereby this is clearly separated from information which is not relevant for data protection (e.g. general terms and conditions). Online, it may be an idea to consider working with a so-called 'layered' privacy statement. This is a privacy statement where the first layer contains rather general information about how personal data are processed, and which allows data subjects to click on further layers where they can consult more detailed information, without having to scroll through a long text document.

Article 13 and 14 GDPR

Articles 13 and 14 specify in more detail when what information needs to be communicated to a data subject. Usually, this information is included in a privacy statement published online or distributed as a printed copy.

In the event of direct collection of data from the data subject (Art. 13 GDPR), the following information, inter alia, must be provided:

- the contact details of the data protection officer, where applicable;
- the purposes of the processing for which the personal data are intended as well as the legal basis for the processing;
- the recipients or categories of recipients of the personal data, if any[142];
- the period for which the personal data will be stored, or if that is not possible, the criteria used to determine that period. A data subject must be able to have an idea, taking into account the individual situation, about how long certain personal data will be stored;
- which rights the data subject has, including the rights a data subject has in the event of automated decision-making (i.e. right to human intervention, right to express an opinion and right to challenge the decision) and how these can be exercised;
- the right to lodge a complaint with a supervisory authority.

Under Article 14 (indirect data collection e.g. through a third party[143]) the same information must be communicated, along with: [144]

- which categories of data are processed (in this case, the data subject is not aware of what data have been collected about him/her);
- from which source the personal data originate, and if applicable, whether it came from publicly accessible sources. In principle, the specific source of the data (surveys, online or mobile applications, open data, social media, links between different data sources, etc.) must be mentioned here, unless this is impossible.[145] In such cases, the nature of the source must be stated (e.g. public or private source) and the type or sector of the source.

With regard to the **moment when this information must be communicated**, the following points should be taken into account:

[142] Recipients are all parties that obtain personal data through the controllers, for example other controllers or processors. In principle, these must be listed individually, but it is also sufficient to list them by category, for example by making clear the sector and location of the recipients.

[143] These could be other controllers such as a previous employer of the data subject, data brokers, other data subjects or public sources including the Official Gazette, CBE, etc.

[144] However, it is good and transparent practice to also communicate this information when data are collected directly from the data subject.

[145] The fact that different information on data subjects in a database come from different sources does not mean that the source cannot be identified, irrespective of the time investment or workload imposed on the controller

- if the data are collected directly from the data subject, the information must be provided at the time of the data collection (for example, if a data subject fills out an online web form, subscribes to a newsletter or uploads a CV);
- If the data are collected indirectly, it depends on the situation:
    - the principle is that the information must be provided within a reasonable time, depending on the actual circumstances in which the personal data are processed, but at the latest within one month after the personal data is obtained;
    - if the personal data are used for communicating with the data subject (for example, to invite someone to apply for a given job or to purchase particular goods), the information must be given at the latest at the time of the first contact with the data subject; or
    - if a disclosure to another recipient is envisaged, at the latest when the personal data are first disclosed.

However, there are some **exceptions to this information obligation in the case of indirect data collection** as provided for in Article 14 of the GDPR. For example, this information does not have to be provided in cases where personal data are processed for scientific or statistical purposes and (i) the provision of such information proves impossible or would involve a disproportionate effort, or (ii) the provision of such information is likely to render impossible or seriously impair the achievement of the objectives of that processing.[146]

Information requirements specific to AI systems

In Articles 12 to 14 and to a lesser extent in Articles 15 and 22, the GDPR imposes various transparency or information obligations that are crucial for AI applications. Schematically, these can be presented as follows:

GDPR PROVISION	**INFORMATION**
Art. 13(2)(f) - 14(2)(g) - 15(1).	Obligation to inform about the **existence and use** of automated (individual) decision-making and profiling
	Obligation to provide useful information on the underlying **logic**
	Obligation to inform about the **importance** and expected **consequences** of this processing for the data subject.
Art. 22 - Recital 71	Obligation to provide **explanations** in the case of an automated individual decision

Two **preliminary remarks** should be made before discussing these provisions of the GDPR in more detail:

- There is a **difference between the information requirements** in **Articles 12-14** and those in **Articles 15 and 22**. In the first case, the information must in principle be provided before or at the time of the processing of personal data. In the second case, the information will usually only be provided after the data subject requests it. The latter is also more in line with what is understood in AI circles under the explainability question. Articles 15 and 22 of the GDPR are discussed below. As mentioned earlier, we will not discuss the explainability issue in further detail. This section is therefore limited to the information which must be communicated in advance.[147]

- These specific information obligations are (in principle) **only applicable** if 'automated decision making' is involved, whether using profiling or not[148], and there are legal consequences for the data subject or

[146] See also section 4.5 on the processing of personal data for scientific or statistical purposes.

[147] It seems relevant to discuss the explainability question (i.e. supplying information after a decision has been taken) separately.

[148] The GDPR defines profiling as any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements. See also section 5.4 on automated decision-making and profiling.

if the data subject is otherwise *significantly* affected. Nevertheless, it is recommended to communicate the information discussed below also in cases where the automated decision-making process does not produce legal or other significant effects. Indeed, by providing sufficient information to the data subject, taking into account the concrete context of the processing, a controller will ensure that the processing is done fairly and transparently.[149]

Finally, it is recommended to use innovative **visual and interactive techniques** when striving for algorithmic transparency since this often involves communicating complex information. Text that is difficult to understand often has a discouraging effect. Controllers may be able to avoid this by using visual aids when communicating information. A **layered** privacy statement can also be a useful instrument. It is recommended to include at least the information discussed under points 1 and 3 below in the first layer and the information under point 2 in a further layer.

The **following information requirements** are specifically relevant for AI systems:[150]

1. Obligation to inform about the existence and use of automated (individual) decision-making and profiling

Where controllers (or the processors they appoint) use automated decision-making in processing personal data, they must **inform** the data subject. Data subjects therefore need to be informed when they interact directly with an AI system or when they communicate personal data to such systems.

Application

For example, this is the case if a web shop uses a chatbot that communicates with visitors and which, if visitors meet certain conditions, can autonomously give a discount or another bonus. Visitors therefore need to be informed at the start of the interaction that the conversation with the chatbot is an automated process without human intervention and that they can find further information in the data protection declaration.

2. Obligation to provide useful information on the underlying logic

The complexity of AI systems can make it challenging to concisely show how an automated decision-making process or profiling works. Consequently, a controller is not expected to provide a complex explanation of the AI system used, and definitely not to disclose or detail an algorithm, the underlying source code or trade secrets.[151] However, it is important to inform a data subject in an **easily understandable, yet adequately specific, useful and meaningful way** about this often complex processing and its underlying logic. [152]

As such, a controller does not have to provide a complex mathematical explanation of how the algorithms or ML used work. Nonetheless, a data subject must be usefully informed in a simple, clear and understandable way so that the basis on which the AI system makes decisions becomes transparent and the data subject knows what results to expect.

This can be done by communicating the **following information**:

- The categories of data/information (and related attributes) that were or will be used in the (re)training, testing or operational use of the profiling or automated decision making systems. These include the

[149] Recital 60 GDPR. However, this does not rule out that in case of an audit or investigation this information must be communicated to e.g. the supervisory authority or legal body.

[150] If there are specific AI audits or certifications in the future, it may be advisable to also communicate when such an audit last took place or if such certification is available.

[151] See also recital 63 of the GDPR.

[152] In this regard, bear in mind that recital 58 of the GDPR states that this is of particular relevance in situations where the proliferation of actors and the technological complexity of practice make it difficult for the data subject to know and understand whether, by whom and for what purpose personal data relating to him or her are being collected.

personal data collected and how it is collected, the data quality or age of the data. It is also recommended to communicate how the necessary measures were taken to ensure that the training and test data were (and still are) representative of the target group(s) for whom the aim is to make predictions or decisions.

- Why these categories are considered relevant and their respective weightings.[153]
- How the model/profile used in the automated decision making process is constructed, including any relevant statistics used in the analysis.[154]
- Why this profile is relevant to the automated decision making process or what purpose is intended.
- How the profile is used to make a decision about the data subject and what criteria are used. If possible, a controller can explain various elements, such as the main methodological choices regarding, inter alia, the algorithms and/or model structure used, the way in which any parameters are determined[155] and how they contribute to a decision. In this regard, it is also recommended to communicate the performance or accuracy of the underlying model, as tested on independent and representative test data.
- To what extent human control and/or intervention is (possible) on the processing.

This rather simple information will be **more relevant for the data subject** than the underlying mathematical mechanisms and will therefore contribute to the transparency of the processing. Conversely, if data subjects or supervisors were given only a limited and complex explanation of the logic behind such decisions, it is very difficult for them to recognise suboptimal (or erroneous) decisions or unreliable models/systems, whereby they would effectively be prevented them from exercising their rights or powers.[156] Nevertheless, the information listed above is information that will typically not be found in the first layers of a layered data protection declaration.[157]

It is also important to know that this applies to both '**static' AI systems and 'learning' AI systems** where the underlying logic can change over time. In the latter case, the data protection declaration will therefore need to be updated regularly.

APPLICATION

In recruitment, for example, an AI system can assign a score to CVs, indicating the relevance of a particular CV to a vacancy. The controller must then make clear the logic behind this score. For example, it can be explained that this process helps them make fair and responsible recruitment decisions, and the information listed above can be communicated. In addition, the controller can also communicate that the CV scoring methods used are regularly tested to ensure that they remain fair, effective and unbiased.

3. OBLIGATION TO INFORM ABOUT THE IMPORTANCE AND EXPECTED CONSEQUENCES OF THIS PROCESSING FOR THE DATA SUBJECT

This obligation means that a controller must inform the data subject with regard to the **intended or expected processing** and the **way in which** the automated decision-making can influence the data subject/what consequences this may have for him/her. In order for this information to be meaningful and understandable, it is recommended providing real, tangible examples of the type of possible effects.

[153] Including which variables or features were considered for inclusion in the model and which variables/features were selected for the final model.

[154] For example, this relates to the definition and specifying of class labels, any pre-processing applied, the way in which meaningful points in the training and test set are determined from the total population about whom the aim is to make predictions, uncertainties, the training methodologies used and the frequency of re-training, etc.

[155] Such as, for example, the cut-off of the model.

[156]See also section 5.3 on the rights of data subjects. In individual cases and in the context of Article 22 of the GDPR, one way to fulfil this would be by providing a counterfactual explanation that identifies the factors that would need to change in order to achieve a different outcome.

[157] However, the Spanish DPA considers that this information should also be included in the first layer of a layered data protection declaration.

APPLICATION

To illustrate the importance and intended consequences of the processing in the e-commerce case, a controller may explain that visiting different online stores that sell the same type of goods, or just a variety of goods, will result in ads that will offer a certain type of goods/a variety of goods respectively. This can be further illustrated by an instructional video comparing fictional people browsing the Internet, with different browsing habits. It can also be explained that diversified browsing habits will lead to a more complete personal profile being built up, which in turn could result in a greater variety of personalised ads, but also for example discounts.

C. Explained in brief: explainability

In view of the above information, it is recommended that AI developers try to adopt an explainability by design approach and pursue the auditability of systems, ensuring (algorithmic) transparency through the design of an AI system and making it possible to rapidly reconstruct decisions. This should clearly be considered if the AI systems in question have a potentially negative impact on the fundamental rights of individuals.

If such approaches are not possible or desirable, AI developers should have other methods at hand (such as reverse engineering) to **ensure that they can somehow extract an explanation from the AI system** after a **given outcome has occurred** (post-hoc interpretability).[158] Indeed, being able to explain an AI system's decision is inextricably linked to how reliable people will deem the technology to be.

In addition, **internal steps** can be taken to enhance the transparency of AI systems. For example, documentation may be drawn up relating to the AI models developed or used and containing the following information: the intended applications, the technical principles of the model and related parameters, the training datasets and methodologies used, who played what role at what point in the development process and how and when the performance of the AI system was evaluated. A first step in this direction could be the drafting of **data sheets** on the datasets used or **info sheets** on the available AI models. This documentation should ideally be adapted to the intended audience (such as senior management, company lawyers, etc.) and should aim to avoid errors of interpretation or use. A controller that appoints a processor that uses AI systems in its services must ensure that its staff also has access to such information.

5.2. What limitations does the GDPR impose on the storage of personal data?

ESSENCE

The principle of storage limitation means that personal data must be deleted or made anonymous as soon as they are no longer necessary for the purposes for which they were collected. [159]

IN ACTION

✓ Specify storage periods for each type of data, possibly by processing activity. If it is not possible to determine storage periods, specify the parameters whereby the time frame can be determined.

✓ Regularly consider whether the personal data processed is still needed, and if not, delete and/or anonymise it.

✓ For archiving, research or statistical purposes, personal data can be kept longer, but they can only be used further for this purpose.

✓ Bear in mind that individuals have the right to have their data deleted in certain circumstances.[160] Estimate what impact this right might have on the operation, development and roll-out of the AI application and take this into account.

[158] This is the examination of a product (usually software or a communication protocol) to derive the requirements that the product is trying to meet, or to ascertain its exact internal functioning.

[159] Art. 5(e) GDPR.

[160]See also section 5.3 on the rights of data subjects.

✓ This principle has a highly technical character, therefore collaborate with other departments within the organisation.

The GDPR itself **does not impose specific storage periods** for different types of data. It is up to the **controller to determine this** and it will depend on how long the data is needed for the specific processing operation. However, sometimes statutory periods will be relevant, for example accounting deadlines or statutes of limitations.

The **method of storage** can also have an impact on the storage periods. For example, the storage period for locally stored data , such as in a robot or voice assistant, will need to be shorter than for data stored in a central location, because of the differences in storage capacity and security capabilities. The application of storage periods is partly a **technical matter**, which requires **cooperation** between different departments, such as IT, business and compliance.

APPLICATION

Employers collect a lot of data from applicants when they apply for a job (CVs, letters of motivation or recommendation, medical examinations, etc.). Once the recruitment process is completed, these data must in principle be destroyed, as the purpose for which they were processed, i.e. recruitment, has been achieved. If the employer wishes to create a recruitment reserve, he must ask the applicant's permission in this regard. The storage of this reserve must also be limited in time to. If the employer wishes to keep the data of candidates in order to reject future applications from the same persons, this may only be done for a certain period of time (e.g. 2 years) as laid down in a data protection policy.

The general rule is therefore that personal data may **not be stored indefinitely**, nor may they be stored solely because they might be 'useful' in the future. There is an exception for archiving, research or statistical purposes.[161] In the latter case, however, appropriate technical and organisational measures must be taken, such as anonymisation, pseudonymisation or access limitation. If data are kept on this basis indefinitely, they may not be subsequently used for any other purpose.

Deleting data that is no longer needed **reduces the risk that it will be irrelevant, excessive, inaccurate or obsolete**. They can also no longer be the subject of a data breach.

APPLICATION

One possible example relates to rankings of candidacies produced by an AI system. Such rankings lose their relevance (in principle) once the recruitment procedure is completed and should therefore be deleted. (On the other hand, the accuracy score given by an employee to a particular ranking may be kept longer).

Having a robust data retention policy in place also helps adhere to the principles of **data minimisation**[162] **and accuracy**[163]. It also reduces the risk of such data **being used erroneously**, to the detriment of the data subjects.

If an organisation uses a **processor** to process personal data, the processor must of course also comply with the obligations of the GDPR. Upon completion of the processing services, the processor must also erase or return the personal data to the controller. Existing copies must be deleted unless further storage is required by law or is part of the service. The organisation also has the right to have an **audit** carried out of the processor to verify that it is fulfilling its obligations.

[161] See also section 4.5 on the processing of personal data for scientific or statistical purposes.

[162] According to article 5(1)(c) of the GDPR, personal data shall be adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed ('data minimisation'); See also section 4.2

[163] According to article 5(1)(d) of the GDPR, personal data shall be accurate and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay ('accuracy').

Data subjects have the **right to erasure** in certain cases. For example, when the data is no longer necessary for the purpose for which it was collected or when the data was collected on the basis of consent and the data subject wishes to withdraw his/her consent.

APPLICATION

Persons who have applied for a job at a given company and initially gave consent to be included in the recruitment reserve may withdraw their consent at any time. Consequently, the related data must be deleted.

5.3. What rights do data subjects have when their data is processed by AI systems?

ESSENCE

The GDPR attributes natural persons a number of rights that allow them to maintain control over their data. These rights apply to personal data used in all the different stages of development and implementation in an AI system, including personal data (1) included in training data, (2) used to make a prediction during use, and (3) that could be contained in the model itself.

IN ACTION

- ✓ Set up internal procedures to respond to all types of requests from individuals seeking to exercise their rights.
- ✓ Respond to such requests as soon as it is clear that someone wishes to exercise his/her rights.
- ✓ Bear in mind that the deadline for replying to a request is (in principle) one month, but that in certain cases it can be extended by two months. Inform the person of this.
- ✓ Be aware that requests can be refused in specific cases, but that such refusal must be justified by the organisation in question. In addition, it must always be communicated that a person has the right to lodge a complaint with the DPA and has the option of appealing to the (civil) courts.
- ✓ Keep in mind that in principle information should be provided free of charge, but if the request is manifestly unfounded or excessive, a reasonable fee may be charged. The request may also be refused. In this regard, it must be demonstrated that the request is manifestly unfounded or excessive.
- ✓ Verify the identity of the requesting party if necessary, but not in an unreasonable way.

A. Common provisions

The GDPR does not specify how an individual can make a valid request. Therefore, it is important to **outline an (internal) procedure** and, for example, clearly indicate in a privacy statement how individuals can exercise their rights vis-à-vis the organisation (such as how to submit their request: mail, online form, etc.). Once it is clear that an individual wishes to exercise his/her rights, the organisation must **respond within one month**.

If it refuses to handle a request, it must give reasons for the **refusal** with regard to the applicant. This must also be within the month.[164] Where the request is complex or where several requests are made simultaneously, the deadline **may be extended, if necessary, by two months**. The organisation must notify the individual of an extension.[165]

Providing the requested information to the individual is (in principle) **free of charge.**[166]

Where requests from an individual are **manifestly unfounded or excessive**, for example due to their repetitive nature, the organisation may choose between:

- charging a reasonable fee in light of the costs;

[164] Art. 12(4) and Recital 59 GDPR.
[165] Art. 12(3) GDPR.
[166] Art. 12(5) GDPR.

- refusing to act on the request.[167]

It is up to the organisation to prove that the request is manifestly unfounded or excessive.

Where the organisation processes a large quantity of personal data, it must be able to request that, before the information is delivered, the data subject specifies the **information or processing activities** to which his/her request relates.[168]

The organisation shall take all **reasonable measures** to **verify** the **identity** of a person requesting access. The requirement to prove identity must be reasonable, not disproportionate and not be used to delay or impede the exercise of the right by the data subject. An organisation cannot retain personal data for the sole purpose of being able to react to potential requests.[169]

In an AI context, the possibility exists for a person to **exercise his/her rights at one of several stages** in the life cycle of an AI system which processes personal data. These different phases are briefly discussed below.

Training

With respect to training data that, for instance, has been converted into another form, it becomes less easy to link it to a specific person. However, this does not automatically mean that the data is non-personal. Even if the data has no specific identifiers or contact details, training data can still be considered personal data. Indeed, these data could be used to distinguish a person, on their own or in combination with other data held by the organisation.

For example, the training data in a purchase prediction model may contain a pattern of purchases unique to one specific customer. Therefore, it is important to also take this data into account when there are requests from individuals who want to exercise their rights under the GDPR.

Output

The output of an AI system can be stored in a profile of a person and used to take certain actions related to that person. For example, the product range that a customer sees on a website may have been prompted by the output of the AI system integrated with his/her profile and on which the system makes predictions. Where such profile data is personal data, it must be subject to the rights of access, rectification and erasure.

Where individual inaccuracies in training data may only have a negligible effect on the result, an inaccurate output from a model can directly affect an individual. An error in the recruitment data by the organisation will presumably not have an impact on the training of the model, but it can have a major impact on the person, for example because an incorrect degree is linked to the person.

Model

Sometimes a model may contain a set of individual examples that are part of the internal logic. This is done so that the AI system can distinguish with or between new examples during operationalisation.

Despite the fact that such a model contains only a small percentage of such examples, there is still a chance that a person may wish to exercise his/her rights. Therefore, it is important that such models allow for the easy retrieval of these training examples so that such requests can be followed up quickly.

Such requests can have a **major impact** on an AI system. In the case of the right of access, little or nothing will have to change in the model. However, if it is a request for rectification or erasure, there is a small chance that the model will probably have to be retrained or even destroyed, for example if the personal data processed are inseparable from the model.

[167] Art. 12(5) GDPR.
[168] Recital 63 GDPR.
[169] Recital 64 GDPR.

There is also the possibility that personal data are **disclosed 'by accident'**. In such cases, third parties may be able to access certain elements of the training data or infer who is in this training data by analysing how the model behaves. It will then be difficult to respond to requests from individuals. As such, it is recommended to **regularly and proactively evaluate** whether personal data can be inferred from the models in order to minimise the risk of accidental disclosure.

B. Right of access

ESSENCE

The right of access gives a person the right to obtain information about his/her processed personal data and to receive a copy of it. Among other things, a person has the right to know for what purposes the data are being processed, what personal data are involved and to whom the data may be sent. The right must be exercised in a simple manner. The information provided must be concise, transparent and intelligible. It must be conveyed in an easily accessible form and in clear and plain language.

IN ACTION

- ✓ Check whether personal data is present in the training data and the model. If so, verify that this personal data can be easily retrieved.
- ✓ Bear in mind that requests for access to training records cannot be considered manifestly unfounded or excessive merely because it is more difficult to comply with.
- ✓ Keep in mind that, conversely, it is not necessary to collect or retain additional information just to allow the organisation to identify individuals within training data, for the sole purpose of acting in compliance with the GDPR. Consequently, the organisation may not be able to identify the individual in the training data (and the individual may not be able to provide additional information that would make it possible to identify them), and thus the organisation may not be able to fulfil a request for access. The organisation will inform the data subjects.
- ✓ Keep in mind that a person also has the right to inspect so-called derived data, for example the profile that you have drawn up about him/her.

The right of access gives a person the right to **obtain information** about his/her processed personal data and to receive a **copy** of it.[170] A person has the right to receive information about whether or not the organisation is processing his/her personal data. If this is the case, a person is entitled to the **following information**:

- the purposes of the processing;
- the recipients or categories of recipients to whom the personal data have been or will be disclosed, in particular recipients in third countries or international organisations;
- where possible, the envisaged period for which the personal data will be stored, or, if not possible, the criteria used to determine that period;
- the existence of the right to request from the controller rectification or erasure of personal data or restriction of processing of personal data concerning the data subject or to object to such processing;
- the right to lodge a complaint with a supervisory authority;
- where the personal data are not collected from the data person, any available information as to their source;
- The existence of automated decision-making, including profiling[171];
- the protection measures applicable if personal data are transferred to a third country or an international organisation[172].

[170] Art. 15(1) GDPR.

[171] See also section 5.4 on profiling.

[172] Art. 15(2) GDPR.

Access must also be granted to **derived data**. These are data concerning a person that has been generated by the organisation itself, for example through data analysis. This may be a profile drawn up by the organisation in the context of recruitment.

The information provided must be **concise, transparent and intelligible**. It must be drawn up in **an easily accessible form** and in **clear and plain language**.

Moreover, it must be possible to exercise this right in a **simple manner**.[173] It is therefore advisable to establish a **digital procedure**. Indeed, the GDPR allows that if a person makes an electronic request, the information can be communicated electronically unless the person requests otherwise.[174]

The GDPR acknowledges that the right of access may have a negative impact on the rights of others by stating that this right **must not affect the rights and freedoms of others**. Recital 63 states that this could be extended to trade secrets or intellectual property. However, it must not lead to a refusal of access to all information.

C. Right to rectification

ESSENCE
Individuals have the right to have inaccurate personal data rectified and incomplete data completed.
IN ACTION
✓ Keep in mind that the steps and procedures that are required, will depend on the nature of the personal data and the purpose for which they are used. The greater the importance of the data for training an AI system, the greater the effort must be to verify that this data is correct and, if necessary, to adapt it. ✓ Realise that the right to correct inaccurate data may include training data used for AI systems. ✓ In practice, a request for rectification cannot be refused because the organisation believes that the data used have less impact on the ultimate purposes. ✓ Inform the person if the personal data are nevertheless correct. In this regard, state the reasons why the rectification was refused.

Individuals have the right to have **inaccurate personal data corrected** and **incomplete data completed**.

This right is closely linked to the **principle of accuracy**.[175] Although an organisation must normally have taken the necessary steps with respect to accuracy at the time it obtains the personal data, this right implies a specific obligation to review the accuracy when requested by the data subject.

If the organisation has transferred data to a third party, it must **inform** the third party of the request, unless this proves impossible or involves disproportionate effort.[176]

The organisation must take the necessary steps to **verify that the personal data is correct** and **update** it if necessary. In this regard, the arguments and evidence given by the person must be taken into account.

As previously stated, which steps are required will depend on the nature of the personal data and the purpose for which they are used. The **more important the data are for training** an AI system, the **greater the effort** must be to verify that it is correct and/or must be adapted.

If an organisation decides that the personal data **are correct**, the requesting party must be informed. In doing so, the organisation should also motivate the refusal.[177]

[173] Recital 63 GDPR.
[174] Art. 15(3) GDPR.
[175] Art. 5(1)(d) GDPR.
[176] Art. 19 GDPR.
[177] Art. 12(4) GDPR.

D. Right to erasure ('right to be forgotten')

ESSENCE
If an organisation no longer has a good reason to process personal data, the data must be deleted.
IN ACTION
✓ Pseudonymise or anonymise personal data as soon as possible, or as soon as they no longer need to be processed. ✓ Take all requests for data erasure into account. However, this is not an absolute right, as the GDPR provides for several exceptions.

The **right to data erasure** means that an individual has the right to have his/her personal data erased by the organisation that processes the personal data.[178] However, this right is **not absolute** and only applies in certain circumstances such as when:[179]

- the personal data are no longer necessary in relation to the purposes for which they were collected/processed;
- the person withdraws his/her consent and there are no other legal grounds for the processing;
- the person objects to the processing;
- the personal data have been unlawfully processed;
- the organisation is legally obliged to delete the data after a certain period of time;
- the personal data has been collected from children.

APPLICATION

For example, if training data is no longer needed because the model has already been trained, the organisation must comply with any requests. In some cases, such as when development of the AI system is still in progress, it may be necessary to retain certain training data for the purpose of re-training, quality research, or refining and evaluating the AI system. In such situations, the organisation must consider, on a case-by-case basis, whether it can meet the request.

A request for data erasure does not have to **be granted** in the following cases:

- if the processing is necessary for the exercise of the right to freedom of expression and information (for example, in a newspaper article);
- if the organisation is legally obliged to process the data or to fulfil a task of general interest (e.g. tax data);
- if the processing is necessary for reasons of public interest in the area of public health;
- if the processing is necessary for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes (e.g. national archives).
- if the processing is necessary for the establishment, exercise or defence of legal claims.[180]

APPLICATION

These exceptions are presumably not relevant when training the AI systems discussed in the applications in this guide. However, they may be applicable to other AI systems and therefore must be considered on a case-by-case basis. The exercise of this right could have a major impact on the development of AI systems in general. Indeed, deleting personal data from a large data set may affect the accuracy and reliability of the system.

[178] Art. 17 GDPR.
[179] Art. 17(1) GDPR.
[180] Art. 17(3) GDPR.

It is therefore important to consider pseudonymising or anonymising data before it is used. Moreover, consideration must be given to the extent to which the AI system and the personal data it processes are or can be separated.

E. Right to restriction of processing

ESSENCE

In certain circumstances, individuals may request that an organisation cease actively processing/using their personal data, without deleting the data.

IN ACTION

- ✓ Realise that processing personal data covers a wide range of operations such as collection, structuring, dissemination and deletion of data.
- ✓ In this regard, one of the following options can be used:
  - o temporarily transfer the selected personal data to another processing system;
  - o make the selected data unavailable to users;
  - o temporarily remove published data from a website.

In certain circumstances, individuals may request that the **processing of their personal data be restricted.** They can exercise this right in the following cases:

- When a person disputes the accuracy of the personal data and the organisation verifies this during a certain period;
- When the processing is unlawful and the person opposes the erasure of the personal data and requests the restriction of their use instead;
- When the organisation no longer needs the personal data for the processing purposes, but the individual needs it in the context of legal proceedings;
- When the person has objected to the processing, pending the answer to the question whether the legitimate interests of the organisation outweigh those of the person.[181]

It is important that organisations implement (technical) **procedures** that enable them to limit the processing of personal data if necessary. They should also make sure that the tools they purchase make this possible. In this regard, it should be kept in mind that processing includes a wide range of operations such as collection, structuring, dissemination and erasure of data.

In many cases the restriction will **be temporary** and can therefore be lifted. However, the persons must be informed before the restriction on processing is lifted.[182]

F. Right to data portability

ESSENCE

Individuals have the right to obtain their personal data in a structured, commonly used and machine-readable format. In certain circumstances, an individual may also request an organisation to transfer their data.

IN ACTION

- ✓ Be aware that in many cases the data used to train an AI system (e.g. purchasing behaviour) is provided by the person themselves. This right will then apply if the processing is based on consent or on a contractual ground.
- ✓ Bear in mind that the data may have been transformed in such a way that it can be more easily analysed by the algorithm. If this transformation is significant, the data may no longer be considered to have been

[181] Art. 18(1) GDPR.
[182] Art. 18(3) GDPR.

'provided by the individual' (derived data). In such cases, it is not possible to exercise this right. The other discussed rights, such as access and rectification, nevertheless persist.

✓ Bear in mind that the initial form of the data remain subject to the right of portability.

The right to data portability gives individuals the right to obtain the personal data they have provided in a **structured, commonly used and machine-readable form** from the organisation they provided it to. In addition, they have the right to transfer that data to another organisation without being obstructed by the organisation to which the personal data was originally provided.[183] However, this right **only applies** if:

- the legal grounds for the processing are the consent of the individual or the necessity for the performance of a contract; and
- the processing is carried out by automated processes.[184]

The individual also has the right to have the personal data **transferred** directly from one organisation to another, if this is technically possible.[185] The right to transfer data **applies** only if these data:

- are personal data concerning the person in question; and
- were provided by the individual to an organisation.

In many cases, these personal data are relatively easy to identify (for example, their name, email address, telephone number or age). However, the notion of 'provided' is **broader** than just these cases. It also covers data resulting from **observation of the person's activities**.

According to the Working Party 29, 'provided' refers to personal data that **can be inferred from users' activities**, such as raw data processed by a smart meter or other types of connected devices, activity logs, internet usage history or search queries. What this does not include, is data created by the organisation.[186]

In summary, the **following categories** qualify as 'provided by the person':

- Data actively and knowingly provided by the individual (such as email address, user name and age);
- Observation data provided by the individual through the use of a service or device. This can include an individual's search history, internet traffic, behaviour on a website and location data. This may be relevant to the AI systems used in the applications discussed in this guide.

On the other hand, it does not include data that the organisation **derives and reduces** on the basis of this provided data.

Application

For example, the profile that an AI system creates for the organisation in the context of job applications (e.g. to assess a person's suitability for a given job) cannot in itself be considered as 'provided by the person'. Another example is when a web shop allows customers to download their sales history (observation data), but not the recommendations that an AI system derives based on this so that customers are shown products that they may find interesting (derived data).

Finally, the right to data portability **may not infringe the rights and freedoms of third parties** (for example, in the context of professional secrecy or privacy).[187]

[183] Art. 20(1) GDPR.
[184] Art. 20(1) GDPR.
[185] Art. 20(2) GDPR.
[186] This can be based on the observed data or received directly through input, for example a user profile created based on an analysis of the collected raw data about historical purchasing behaviour or viewed job vacancies.
[187] These are other persons and data covered by intellectual property and trade secrets.

G. Right to object

ESSENCE

In certain circumstances, an individual may ask an organisation to stop processing personal data because of the specific situation or in the case of direct marketing purposes, including profiling.

IN ACTION

- ✓ Keep in mind that with direct marketing, the data subject always has the right to object without a reason. As a result, the organisation must automatically cease processing for this purpose.
- ✓ Bring the possibility of exercising the right of objection to the attention of the data subject clearly and separately from other information.

This provision is aimed at processing operations which have **valid legal grounds** but which are performed **against the will of a person**. The right can be invoked in **three situations**:

1. THE SPECIFIC SITUATION OF A PERSON

An individual has the right at any time, in specific circumstances, to object to the processing of their personal data based on any of the following legal grounds:

- the protection of the legitimate interests of the organisation;
- the necessity for the performance of a task carried out in the public interest or in the exercise of official authority vested in the organisation, including profiling based on these provisions.[188]

As a result of this specific situation of an individual, his or her interests override those which serve as legal grounds for the processing. This allows the person to object to the processing in question. The specific situation of the individual may be based on their rights or freedoms such as family circumstances or professional confidentiality, such as lawyers.

The organisation must therefore cease processing the personal data unless the organisation can demonstrate that:

- there are compelling legitimate grounds for the processing which override the interests, rights and freedoms of the individual; or
- there are grounds relating to the establishment, exercise or defence of legal claims.[189]

It is up to the organisation to demonstrate that the compelling legitimate interests **override** the interests or the fundamental rights and freedoms of the individual.[190]

2. DIRECT MARKETING

A person has the right to object, at any time and without stating reasons, to the processing of his or her personal data processed for the purposes of direct marketing, including profiling used for that purpose.[191] For example, a web shop that uses an AI system that sends personalised advertising based on customers' buying behaviour.

Unlike the previous ground, the organisation cannot contest this. It is an absolute right and the processing of personal data must be stopped.

3. SCIENTIFIC OR HISTORICAL RESEARCH OR STATISTICAL PURPOSES

[188] Art. 21(1) GDPR.
[189] Art. 21(1) GDPR.
[190] Recital 69 GDPR.
[191] Art. 21(2) GDPR.

Where personal data are processed for scientific or historical research purposes or for statistical purposes, a person, on grounds relating to his or her particular situation, has the right to object to the processing of personal data concerning him or her, unless the processing is necessary for the performance of a task carried out for reasons of public interest.[192]

5.4. What does the GDPR state about automated individual decision making and profiling and how does this impact AI systems?

ESSENCE

AI systems are only in certain cases allowed to make fully automated decisions, i.e. without human intervention, whereby these decisions have legal or similar consequences for individuals. AI systems that purely support or enhance human decisions are not covered by this restriction. The human intervention must be meaningful. Situations where a human intervenes in the system only *pro forma* still fall under the stricter conditions.[193] The degree and quality of human assessment and intervention prior to the final decision about a person is an important factor in determining whether or not an AI system makes entirely automated decisions.[194]

Profiling is any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.

AI systems often use the processing of personal data to make or support a decision. This can be the case, for example, by supporting HR teams in identifying candidates for interviews by ranking the applications.

Automated decision-making has a different scope and may partially overlap or be the result of profiling. Exclusively automated decision-making involves making decisions by technological means without any human intervention. In principle, there is a general prohibition on exclusively automated decision-making which produces legal effects concerning an individual or similarly significantly affects such individual. There are three exceptions to this prohibition, namely if the decision:

- is necessary for entering into, or performance of, a contract;
- is authorised by law;
- is based on the data subject's explicit consent.[195]

Automated decision making whereby special categories of personal data[196] are processed is only allowed if one of the above exceptions appliesand the person involved has given his or her explicit permission or the processing is necessary for substantial public interests.

IN ACTION

✓ Carry out a DPIA[197].

✓ Inform individuals that their data will be used for this type of processing. This applies to data collected directly from the individual or from other sources.

✓ Provide information on the underlying logic of the AI system and its potential impact on individuals.

[192] Art. 21(6) GDPR. See also section 4.5 on the processing of personal data for scientific or statistical purposes.

[193] See also: R. Binns and V. Gallo, “Automated Decision Making: the role of meaningful human reviews”, https://ico.org.uk/about-the-ico/news-and-events/ai-blog-automated-decision-making-the-role-of-meaningful-human-reviews/.

[194] See also: R. Binns & V. Gallo, "Automated Decision Making: the role of meaningful human reviews".

[195] Art. 22(2) GDPR.

[196]See also section 3.2 on special categories of personal data

[197] See also section 4.4 on the DPIA.

- ✓ Inform individuals about why automated processing is used and what the possible results and significant consequences are, especially since such processing is not visible to individuals.[198]
- ✓ Establish a policy so that individuals can challenge a decision and explain what the conditions are for doing so.
- ✓ Make sure an evaluation is performed by someone qualified to potentially review the decision.
- ✓ Ensure that the procedure is simple and user-friendly for persons wishing to exercise their rights.
- ✓ Take the following steps in order to take the rights of individuals into consideration:
    - Identify what system requirements are needed to enable a meaningful human assessment from the design phase onward.
    - Organise the necessary and appropriate training/education for employees who oversee the system.
    - Empower and support staff to take into account the concerns of individuals, and if necessary, to override the decision of the AI system.
- ✓ Take the following appropriate measures
    - Conduct regular quality checks of the systems to ensure that individuals are treated fairly and not discriminated against on the basis of special categories of personal data or otherwise.
    - Evaluate algorithms and test the AI systems in order to verify that they actually function as intended and do not generate discriminatory, incorrect or unjustified results.
    - Have independent third parties perform evaluations , especially when decision-making based on profiling may significantly affect individuals. Provide this independent third party with all required information on how the algorithm or ML system works.
    - Obtain contractual safeguards for algorithms developed by third parties that demonstrate that the necessary evaluations and tests have been performed and the algorithm meets the agreed and legal requirements.
    - Take specific measures to ensure data minimisation, apply clear storage periods for profiles and for personal data used for the compilation or application of the profiles.
    - Use techniques for anonymisation and pseudonymisation in the context of profiling.
    - Provide ways for the individual to express his/her stance and challenge decisions.
    - Provide a mechanism for human intervention in specific cases, for example, by providing a link to an appeal process at the time the automated decision is communicated to the individual, with agreed time limits for review and a designated point of contact for questions.

A. Advantages and disadvantages of profiling and automated decision-making

Profiling and automated decision-making can be **very useful** for organisations and also benefit individuals in many areas, especially in industries such as healthcare, education, financial services and marketing. They can lead to faster and more consistent decisions, especially decisions that require a large amount of data to be analysed and decisions to be made quickly.

While these techniques can be useful, there are **some risks** involved:

- Individuals are often unaware of profiling and probably do not expect their personal data to be used in this way.
- Individuals do not understand how these processes work and what influence these processes can have.
- The decisions made may entail major disadvantages for some individuals. There are several examples that demonstrate the disadvantages of using AI, for example, in recruitment. For instance, an algorithm developed by Amazon to scan cover letters was found to disadvantage women.

[198] See also section 5.1. on transparency.

Just because the analysis of data may show a correlation does not mean that it is significant or even relevant. Given that the process can only make assumptions about an individual's behaviour or characteristics, there is always a margin of error which needs to weighed against the risk of actually using the results.

B. Profiling

Profiling is the **automated processing of personal data evaluating the personal aspects relating to a natural person**, in particular to analyse or predict aspects concerning the data subject's performance at work, economic situation, health, personal preferences or interests, reliability or behaviour, location or movements.

Profiling can use algorithms and ML to find correlations between different data sets. Organisations typically use profiling to learn about personal preferences, predict behaviour, and make decisions about individuals.

Profiling therefore consists of **three elements**:

1.	It must be an automated form of processing.
2.	It must relate to personal data.
3.	The purpose of profiling must be to evaluate personal aspects of a natural person.

An organisation **engages in profiling when** it:

- collects and analyses personal data on a large scale and uses algorithms or ML in the process;
- identifies associations to make connections between different behaviours and characteristics;
- creates profiles to apply to individuals;
- predicts the behaviour of individuals based on their assigned profile.

Organisations obtain personal data **from a variety of sources** such as search queries, purchasing habits, or behavioural data via smartphones, social media, and wearable devices. This information is then analysed to classify individuals into different groups or sectors. This analysis can also identify links between different behaviours and characteristics to create profiles for individuals. This profile may subsequently consist of new (derived) personal data about this person.

An example makes this clear. A web shop can collect data on customers such as purchases, products they almost bought, location data or preference for a certain colour. An AI system is used to build profiles and categorise customers into certain segments based on their characteristics (such as age, gender, spending, and preferences). Based on this, the content of the web shop can be personalised, for example by what type of products the customers will see first.

There are **a number of possible ways** that profiling can be used:

- general/regular profiling;
- decision-making based on profiling; and
- entirely automated decision-making, including when based on profiling, which produces legal effects or otherwise significantly affects the data subject.

APPLICATION

The difference between decision-making based on profiling and entirely automated decision making can be explained by the following two examples, where an individual applies for a job online. On the one hand, a person can decide whether to assign a position based on a profile created solely by automated means (decision-making based on profiling). On the other hand, an algorithm may decide whether to grant the position and the decision is automatically disclosed to the individual without prior and meaningful human review (entirely automated decision making, including when based on profiling).

Organisations can use profiling and automated decision-making as long as they comply with **all the principles in the GDPR** and have **legal grounds** for the processing. Additional safeguards and restrictions apply in the case of entirely automated decision making, including when based on profiling.

C. Automated decision-making: the person decides

Automated decision-making has a different scope and may partially overlap or be the result of profiling.

Automated decisions can be based on **any type of data**, for example:

- information **directly** provided by the data subjects, such as questionnaire responses;
- data on the **data subjects that are recorded**, such as location data collected through an application;
- **derived data**, such as a profile of the person that has already been created (for example, a profile on a web shop).

Automated decisions **can be made with or without the use of profiling**. Conversely, profiling can take place without automated decisions being made. Profiling and automated decision-making are **not necessarily separate activities**. Something that starts as a simple automated decision-making process can evolve into automated decision-making based on profiling, depending on how the data is used.

A fully automated system can provide recommendations on individuals. If there is still human intervention, with other factors being considered before a final decision is made, it is not based solely on automated processing.

D. Entirely automated decision-making: the AI system decides

GENERAL PROHIBITION

Entirely automated decision-making entails making decisions by technological means **without any human intervention**. The GDPR provides for a **general prohibition** on decision-making based solely on automated processing, including profiling. This prohibition applies regardless of whether or not the data subject takes any action in relation to the processing of their personal data.[199]

The prohibition **only applies** in the case of:

- decisions based solely on automated processing;
- that have a legal effect on a person or otherwise significantly affect the person.

With entirely automated decisions, there is no **human intervention** in the decision-making process. If such process makes a recommendation in relation to a person, but a human evaluates that recommendation and takes other factors into account in making the final decision, this is not a decision 'based solely on automated processing'.

This prohibition **cannot be circumvented by limited human intervention**. For example, if someone routinely applies automatically generated profiles to individuals without actually affecting the outcome, this is still a decision based solely on automated processing.

[199] Article 22 GDPR.

As mentioned above, the decision must also **have legal effects** or **affect the data subject significantly in some other way**. The GDPR does not define these terms, but it can be inferred that only serious, significant effects are referred to.

A decision that has a legal effect can, on the one hand, **affect someone's (fundamental) legal rights** (such as freedom of association, the right to vote and the right to appeal) and, on the other hand, affect **someone's legal status**. Examples of such an effect are automated decisions about a person that lead to, for example, the termination of a contract or the right to or the refusal of a certain legally awarded social benefit including child benefit or rent allowance.

APPLICATION

A entirely automated decision to no longer include someone in a recruitment process will, in principle, fall under the prohibition (see exceptions, however).

A decision that significantly affects a person has **similar effects** on a person's situation, behaviour, or choices.

APPLICATION

Typical examples are automatic refusal of online recruitment processes and/or processing of applications via the internet without human intervention.[200] Another relevant example is when a company decides to interview certain individuals based on the results of an online aptitude test. This decision has a significant effect, as it determines whether or not a person is eligible for the position.

If there is **uncertainty** as to whether a decision will significantly affect a person in another way, it is important to consider to what extent there are **consequences** for:

- financial circumstances;
- health;
- reputation;
- behaviour; or
- freedom of choice.

APPLICATION

Although **online advertising** that uses automated tools and decision-making does not at first sight fall within the scope of article 22 GDPR, some nuance is nevertheless required.
An advertisement for a common online fashion outlet based on a simple demographic profile such as, for example, 'women aged 25 to 35 in the Brussels region who are likely to be interested in fashion and certain clothing' will not affect a person significantly. However, the decision may significantly affect individuals in other ways, depending on the specific characteristics of the particular case:

- the intrusive nature of the profiling process, such as tracking across different websites, devices and services;
- the expectations and wishes of the data subjects;
- the way in which the advertisement is presented; or
- the use of knowledge regarding the vulnerabilities of the data subjects approached.

EXCLUSIONS

[200] Recital 71 GDPR.

There are **three exceptions** to the general prohibition on entirely automated individual decision-making, namely if the decision:

- is necessary for entering into, or performance of, a contract between the data subject and a data controller;;
- is authorised by law;
- is based on the data subject's explicit consent.[201]

1. Necessary for entering into, or performance of, a contract between the data subject and a data controller

The organisation must be able to demonstrate whether the entirely automated processing is necessary for the achievement of its purpose, namely the conclusion of a contract with the natural person. The organisation must consider whether there is any other method of achieving the objective that is less detrimental to the data subjects in question. If there are other means which are equally effective and less detrimental to the data subject, the processing is not necessary.

Take the example of a company that posts a job vacancy online. Given the popularity of the employer, the company receives thousands of applications. Due to the exceptionally high number of applications, the company finds it virtually impossible to select suitable candidates without first using fully automated means to filter out unsuitable candidates. In this case, automated decision-making may be necessary in order to make a pre-selection of potential candidates, with the intention of ultimately concluding a contract with a data subject.

2. Authorised by law

Automated processing, including profiling, can in principle take place if a law authorises its application. The law must also provide for appropriate measures to protect the rights and freedoms and legitimate interests of the data subject.[202]

3. Explicit consent

Explicit consent is not defined in the GDPR. 'Regular' consent requires 'a statement or by a clear affirmative action'.[203] It is therefore appropriate to clarify what additional efforts an organisation has taken to obtain this explicit consent.

The notion explicitly refers to the way in which consent is expressed by a person and therefore means that a person must provide an explicit statement of consent for that particular processing. An obvious method is a written statement, which should also be signed by the data subject. In a digital context, this can be done for example by filling out an online form, sending an e-mail, scanning a document with a signature or using an electronic signature.

A so-called two-stage verification of consent is also a possibility to render a consent explicit. For example, an e-mail informing the individual of a web shop's intention to process certain personal data. The web shop explains in the email that it is asking for permission to use a specific dataset for a specific purpose. If the person consents to the use of these data, the web shop will ask for a reply by e-mail with the statement 'I agree'. After the response is sent, the person receives a verification link that must be clicked or an SMS message with a verification code to confirm their consent.

E. Special categories of data and children

Special categories

[201] Art. 22(2) GDPR.
[202] Recital 71 GDPR.
[203] Art. 4(11) GDPR.

Automated decision making whereby special categories of personal data are processed is only allowed if one of the above exceptions applies and the person has given his or her explicit permission or the processing is necessary for substantial public interests.[204] The organisation must also take appropriate measures to safeguard the rights and freedoms and legitimate interest of individuals.

It is important to be aware that bringing together different types of personal data can reveal sensitive information about individuals. For example, one study combined likes from Facebook with a simple survey, and was able to predict the sexual orientation of men with 88% accuracy. Moreover, it predicted ethnicity with 95% accuracy and whether a user was Christian or Muslim with 82% accuracy. Such surveys are subject to the same legal obligations under the GDPR as if sensitive personal data had been processed from the outset.

CHILDREN

The relevant provisions in the GDPR **do not distinguish between adults and children**. However, Recital 71 states that entirely automated decision-making, including when based on profiling, which produces legal effects or otherwise significantly affects the data subject, **should not concern a child**.

But since this wording is not found in Article 22 itself, the WP29 did not consider it to be an absolute prohibition with regard to children. In light of this consideration, the WP29 does recommend that, in principle, organisations should not use the exceptions to justify this type of processing.

However, in **exceptional cases**, it may be necessary to use automated decision making only, for example to protect the well-being of children. In such cases, the processing may be carried out under one of the three exceptions discussed. Nonetheless, **appropriate protective measures** suitable for children must be taken. The organisation must ensure that these measures effectively protect the rights and freedoms and legitimate interests of the children whose data they process.

The need for specific protection for children should, in particular, apply to the use of personal data of children for the purposes of marketing or creating personality- or user profiles and the collection of personal data with regard to children in the context of services offered directly to a child.[205]

Organisations should therefore **refrain from applying profiling to children for marketing purposes**. Children are particularly vulnerable in the online environment and are easier to influence with advertising that is tailored to their browsing behaviour. For example, in online games, profiling can be used to show advertising to players whom the algorithm believes are more likely to spend money on the game, and to show more personalised ads. On account of their age and immaturity, children do not always understand the reasons behind this kind of marketing and the consequences it can have for them.

F. Related rights of data subjects

In relation to the right of access, individuals have the **right to receive information** about the solely automated decision-making process used, including profiling. This information relates to:[206]

- the existence of automated decision-making, including profiling;
- useful information about the underlying logic; and
- the importance and expected consequences of this processing for the data subject.

In this context, the organisation must provide the individual with information about the **expected consequences of the processing** and an **explanation of a specific decision**. Every data subject therefore has the right to be

[204] Art. 22(4) GDPR.

[205] Recital 39 GDPR.

[206] Art. 15 GDPR. See also section 5.3 on the rights of data subjects.

informed about how automated decision-making works, including the logic involved and the consequences of such processing, at least when the processing is based on profiling.[207]

Furthermore, an organisation using automated individual decision-making must also provide individuals with the following options/rights:

- the right to human intervention by the organisation that uses automated individual decision-making;
- the right to express their views; and
- the right to contest the decision. It also follows from this last right that individuals must be able to obtain an explanation of the decision taken following such an assessment.[208]

For an individual, these measures include at least a possibility to **obtain human intervention**, **express his/her point of view** and **contest the decision**. Human intervention is crucial in this regard. An assessment of such requests must be performed by someone who is competent and capable of modifying the AI system's decision. This person must (be able to) thoroughly analyse all relevant data, together with any additional information provided by the data subject.[209]

The organisation must **provide individuals with an easy way** to exercise these rights. This once again underscores the need to be **transparent about the processing**. The data subject can only contest a decision or express his/her point of view if he/she fully understands how and on what basis the decision was taken.[210]

6. Conclusion

This exploratory guide has discussed the application of some of the provisions of the GDPR to the design, development and use of AI systems. While not an exhaustive work, we attempted to clarify some of the fundamental provisions of the GDPR in a practical way. We tried to achieve a *multi-layered* approach. Each section started with a summary and an overview of some concrete actions that organisations and users can go through with regard to AI systems. In the coming months, the KCDS will distribute sheets and other practical tools based on the different chapters of this guide. This exploratory guide is therefore by no means an end point, but rather a 'living' document that will be supplemented and refined as necessary. Questions, comments or feedback on/about this guide are welcome and may be addressed directly to the researchers involved.

[207] Recital 63 GDPR. See also section 5.1. on transparency.

[208] This is confirmed in Recital 71 of the GDPR. See also section 5.1. on transparency.

[209] Art. 22(3) and Recital 71 GDPR.

[210] See also section 5.1. on transparency.

7. BIBLIOGRAPHY

- Chapter 2

High-Level Expert Group on Artificial Intelligence. “A definition of Artificial Intelligence: main capabilities and scientific disciplines", 7p., https://ec.europa.eu/digital-single-market/en/news/definition-artificial-intelligence-main-capabilities-and-scientific-disciplines

European Commission, “White paper on artificial intelligence - A European approach to excellence and trust”, 19 February 2020, 30p. https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf.

European Commission, Communication from the Commission to the European Parliament, the European Council, the Council, the European Economic and Social Committee and the Committee of the Regions, “Artificial intelligence for Europe”, 25 April 2018, COM(2018) 237, 23p., https://ec.europa.eu/transparency/regdoc/rep/1/2018/EN/COM-2018-237-F1-EN-MAIN-PART-1.PDF

Knowledge Center Data & Society, Brainfood, "How to Survive a Conversation about Artificial Intelligence," https://data-en-maatschappij.ai/en/publications/brainfood-hoe-overleef-je-een-gesprek-over-artifici%C3%ABle-intelligentie

L. Steels, “Artificiële intelligentie. Naar een vierde industriële revolutie?”, Royal Flemish Academy of Belgium for Sciences and Arts, 2017, 49p., https://www.kvab.be/nl/activiteiten/artifici%C3%ABle-intelligentie-naar-een-vierde-industri%C3%ABle-revolutie.

M.J. Vetzo, J.H. Gerards and R. Nehmelman, *Algoritmes en grondrechten*, Boom Juridisch, The Hague, 2018, p. 244.

R. Nijman, “Cognitive Computing en IBM Watson – Wat is het en wat biedt het de overheid?”, https://www.ibm.com/blogs/think/nl-en/2015/01/12/cognitive-computing-en-ibm-watson-wat-is-het-en-wat-biedt-het-de-overheid/. See also: https://www2.cio.nl/development/85006-wat-is-cognitive-computing.

Regulation 2018/1807 of the European Parliament and of the Council of 14 November 2018, on a framework for the free flow of non-personal data in the European Union, PE/53/2018/REV/1, OJ L 303, 28.11.2018, p. 59-68.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016, on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (GDPR).

- Chapter 3

Commission for the Protection of Privacy, "Een gids om kleine en middelgrote ondernemingen (KMO’s) voor te bereiden op de Algemene Verordening Gegevensbescherming," 32p., www.gegevensbeschermingsautoriteit.be/sites/privacycommission/files/documents/KMO_NL_update.pdf.

G. LaFever, “Anonymisation does not work for big data due to lack of protection for direct & indirect identifiers and easy re-identification vs pseudonymization”, 2019, gdpr.report/news/2019/08/12/anonymisation-does-not-work-for-big-data-due-to-lack-of-protection-for-direct-indirect-identifiers-and-easy-re-identification-vs-pseudonymisation.

Article 29 Working Party, “Opinion 4/2007 on the concept of personal data”, June 2007, 28p., https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2007/wp136_en.pdf.

Article 29 Working Party, “Opinion 5/2014 on anonymisation techniques”, April 2014, 43p., https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf.

Information Commissioner's Office, “Anonymisation: managing data protection risk code of practice”, 2012, 108p., https://ico.org.uk/media/for-organisations/documents/1061/anonymisation-code.pdf.

Information Commissioner's Office, “Guide to the GDPR”, 2019, 317 p., https://ico.org.uk/media/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr-1-0.pdf.

L. Rocher, J.M. Hendrickx and Y. de Montjoye, “Estimating the success of re-identifications in incomplete datasets using generative models”, *Nature Communication*, 2019, vol. 10, no. 3069, https://doi.org/10.1038/s41467-019-10933-3.

M. Finck and F. Pallas, “They who must not be identified – Distinguishing Personal from Non-Personal Data under the GDPR”, Max Planck Institute for Innovation & Competition, Research Paper No. 19-14, 2019, 48 p., https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3462948.

Communication from the Commission to the European Parliament and the Council, “Guidance on the Regulation on a framework for the free flow of non-personal data in the European Union”, COM/2019/250, 29 May 2019, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52019DC0250&from=en.

R.P. Santifor, “Naar een meer genuanceerde benadering van ‘pseudonimisering’ in het privacyrecht”, *P&I,* October 2019, no. 5, 10p.

Regulation 2018/1807 of the European Parliament and of the Council of 14 November 2018, on a framework for the free flow of non-personal data in the European Union, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32018R1807&from=EN.

- Chapter 4

Agencia Española Proteccion Datos (AEPD), “RGPD compliance of processings that embed Artificial Intelligence - An introduction”, February 2020, 49p., https://www.aepd.es/sites/default/files/2020-02/adecuacion-rgpd-ia-en_0.pdf.

Dutch DPA, “Data Protection Impact Assessment”, https://autoriteitpersoonsgegevens.nl/nl/zelf-doen/data-protection-impact-assessment-dpia.

Dutch DPA, ”Toezicht op AI en Algoritmes”, 17 February 2020, 11p., https://autoriteitpersoonsgegevens.nl/nl/nieuws/toezicht-op-algoritmes.

Decision of the General Secretariat No. 1/2019 of January 16, 2019, Belgian Official Gazette March 22, 2019, 28512-28514 www.gegevensbeschermingsautoriteit.be/sites/privacycommission/files/documents/01_2019_AS.pdf.

Centre for Cybersecurity Belgium (CCB), Cybersecurity guide for SME, 2017, 17p., https://ccb.belgium.be/en/document/guide-sme .

Commission for the Protection of Privacy, "Big Data Report," 2017, 54p., https://www.gegevensbeschermingsautoriteit.be/big-data-rapport

European Network and Information Security Agency (ENISA), Big Data Security, Good Practices and Recommendations on the Security of Big Data Systems, December 2015, 30p., https://www.enisa.europa.eu/publications/big-data-security.

European Network and Information Security Agency (ENISA), Big Data Threat Landscape and Good Practice Guide, January 2016, 62p., https://www.enisa.europa.eu/publications/bigdata-threat-landscape.

European Network and Information Security Agency (ENISA), Handbook on Security of Personal Data Processing, December 2017, 68p., https://www.enisa.europa.eu/publications/handbook-on-security-of-personal-data-processing.

European Network and Information Security Agency (ENISA), Privacy by design in big data, An overview of privacy enhancing technologies in the era of big data analytics, December 2015, 80p., https://www.enisa.europa.eu/news/enisa-news/privacy-by-design-in-big-data-an-overview-of-privacy-enhancing-technologies-in-the-era-of-big-data-analytics.

European Data Protection Board, Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, 13 November 2019, p. 9, no. 32-36.

European Data Protection Board, Letter in Response to Sophie In't Veld, Ref: OUT2020-0004, 29 January 2020, 6p., https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_letter_out2020_0004_intveldalgorithms_en.pdf.

European Commission, “White paper on artificial intelligence - A European approach to excellence and trust”, 19 February 2020, 30p. https://ec.europa.eu/info/sites/info/files/commission-white-paper-artificial-intelligence-feb2020_en.pdf.

Article 29 Working Party, “Opinion 5/2014 on anonymisation techniques”, April 2014, 43p., https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf.

Article 29 Working Party, "Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is "likely to result in a high risk" for the purposes of Regulation 2016/679," 248 rev.01, October 2017, 22 p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=611236.

Article 29 Working Party, "Guidelines on transparency under Regulation 2016/679," 11 April 2018, 40p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=622227

Article 29 Working Party, “Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679”, 6 February 2018, 37p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053

Information Commissioner's Office, “Anonymisation: managing data protection risk code of practice”, 2012, 108p., https://ico.org.uk/media/for-organisations/documents/1061/anonymisation-code.pdf.

Information Commissioner's Office, "Big data, artificial intelligence, machine learning and data protection," 4 September, 2017, 114p., https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf.

Information Commissioner's Office, "Data minimisation and privacy-preserving techniques in AI systems", https://ico.org.uk/about-the-ico/news-and-events/ai-blog-data-minimisation-and-privacy-preserving-techniques-in-ai-systems.

Information Commissioner's Office, “Guide to the GDPR”, 2019, 317 p., https://ico.org.uk/media/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr-1-0.pdf.

Information Commissioner's Office, "Guidance on the AI auditing framework. Draft guidance for consultation", 2020, 105p., https://ico.org.uk/media/about-the-ico/consultations/2617219/guidance-on-the-ai-auditing-framework-draft-for-consultation.pdf.

International Conference of Data Protection and Privacy Commissioners (ICDPPC), "Declaration on Ethics and Data Protection in Artificial intelligence," 23 October 2018, 6p. https://edps.europa.eu/sites/edp/files/publication/icdppc-40th_ai-declaration_adopted_en_0.pdf.

K. Wuyts, *Privacy Threats in Software Architectures*, Ph.D., 2015.

Norwegian Data Protection Authority, “Artificial intelligence and privacy”, 2018, 30 p., https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf.

- Chapter 5

Agencia Española Proteccion Datos (AEPD), “RGPD compliance of processings that embed Artificial Intelligence - An introduction”, February 2020, 49p., https://www.aepd.es/sites/default/files/2020-02/adecuacion-rgpd-ia-en_0.pdf.

Article 29 Working Party, "Guidelines on transparency under Regulation 2016/679," 11 April 2018, 40p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=622227

Article 29 Working Party, “Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679”, 6 February 2018, 37p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053

Dutch DPA, ”Toezicht op AI en Algoritmes”, 17 February 2020, 11p., https://autoriteitpersoonsgegevens.nl/nl/nieuws/toezicht-op-algoritmes.

Commission for the Protection of Privacy, "Big Data Report," 2017, 54p., https://www.gegevensbeschermingsautoriteit.be/big-data-rapport

Commission for the Protection of Privacy, "Een gids om kleine en middelgrote ondernemingen (KMO’s) voor te bereiden op de Algemene Verordening Gegevensbescherming," 32p., www.gegevensbeschermingsautoriteit.be/sites/privacycommission/files/documents/KMO_NL_update.pdf.

Conference of the Independent Federal and State Data Protection Supervisory Authorities of Germany (Datenschutzkonferenz, DSK), **"**Hambach Declaration on Artificial Intelligence - seven data protection requirements," 3 April 2019, 4p., https://www.datenschutz-berlin.de/fileadmin/user_upload/pdf/publikationen/DSK/2019/2019-DSK-Hambach_Declaration_AI-en.pdf.

Data Ethics Commission (daten ethik kommission), “Opinion of the Data Ethics Commission – executive summary”, October 2019, 32p., https://www.bmjv.de/SharedDocs/Downloads/DE/Themen/Fokusthemen/Gutachten_DEK_EN.pdf?__blob=publicationFile&v=2.

European Data Protection Board, "Guidelines 05/2020 on consent under Regulation 2016/679," May 2020, 31 p., https://edpb.europa.eu/sites/edpb/files/files/file1/edpb_guidelines_202005_consent_en.pdf.

Article 29 Working Party, "Guidelines on the right to "data portability," WP 242 rev.01, April 2017, 24 p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=611233.

Article 29 Working Party, "Guidelines on automated individual decision-making and profiling for the purposes of Regulation (EU) 2016/679," WP251rev.01, 2018, 47 p., https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=612053.

Data Protection Authority, “Hoe zit dat met de geautomatiseerde beslissingen, waaronder profilering? (Art. 22 GDPR)”, https://www.gegevensbeschermingsautoriteit.be/hoe-zit-dat-met-de-geautomatiseerde-beslissingen-waaronder-profilering-art-22-avg.

Information Commissioner’s Office and The Alan Turing Institute, “Explaining decisions made with AI”, 20 May 2020, 136p., https://ico.org.uk/for-organisations/guide-to-data-protection/key-data-protection-themes/explaining-decisions-made-with-artificial-intelligence/.

Information Commissioner's Office, “Guide to the GDPR”, 2019, 317 p., https://ico.org.uk/media/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr-1-0.pdf.

Information Commissioner's Office, "Big data, artificial intelligence, machine learning and data protection," 2017, 114p., https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf.

Information Commissioner's Office, "Guidance on the AI auditing framework. Draft guidance for consultation", 2020, 105p., https://ico.org.uk/media/about-the-ico/consultations/2617219/guidance-on-the-ai-auditing-framework-draft-for-consultation.pdf.

Information Commissioner's Office, "Automated individual decision-making and profiling", 2018, 23p., https://ico.org.uk/media/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/automated-decision-making-and-profiling-1-1.pdf.

Information Commissioner's Office, "What does the GDPR say about automated decision-making and profiling?", https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/automated-decision-making-and-profiling/what-does-the-gdpr-say-about-automated-decision-making-and-profiling/

International Conference of Data Protection and Privacy Commissioners (ICDPPC), "Declaration on Ethics and Data Protection in Artificial intelligence*,"* 23 October 2018, 6p. https://edps.europa.eu/sites/edp/files/publication/icdppc-40th_ai-declaration_adopted_en_0.pdf.

Joint Research Centre - European Commission (R. Hamon, H. Junklewitz, I. Sanchez), "Technical Report on Robustness and Explainability of Artificial Intelligence – from technical to policy solutions", 2020, 40p., https://publications.jrc.ec.europa.eu/repository/handle/JRC119336

M. Mitchell (*et al*.), "Model Cards for Model Reporting", January 2019, 10p., https://arxiv.org/abs/1810.03993.

Norwegian Data Protection Authority, "Artificial intelligence and privacy", 2018, 30 p., https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf.

R. Binns and V. Gallo, "Automated Decision Making: the role of meaningful human reviews", 2019, https://ico.org.uk/about-the-ico/news-and-events/ai-blog-automated-decision-making-the-role-of-meaningful-human-reviews/.

T. Gebru (*et al*.), "Datasheets for Datasets", March 2020, 24p., https://arxiv.org/abs/1803.09010.