

Artificial Intelligence and Data Protection

How the GDPR Regulates AI

Centre for Information Policy Leadership (CIPL)

March 2020

Contents

I.	Introduction	3
II.	How GDPR Applies Generally in the Context of AI.....	4
A.	Legal Basis	5
B.	Data Protection Principles	5
C.	Accountability	6
D.	Appointing a Data Protection Officer.....	7
E.	Article 30 Inventory.....	8
F.	Controller-Processor Contracts and Data Transfers	8
G.	Data Breaches	9
H.	Individual Rights.....	9
I.	Data Protection by Design and by Default.....	9
J.	Extraterritorial Effect of the GDPR.....	10
K.	Data Protection Authorities, Sanctions and Enforcement.....	10
III.	GDPR Provisions of Particular Relevance to AI	12
A.	Fair Processing	12
B.	Data Minimisation.....	13
C.	Data Protection Impact Assessments (DPIAs).....	14
IV.	GDPR Provisions that Specifically Regulate AI	15
A.	Automated Decision-Making – Definition and Scope.....	15
B.	Automated Decision-Making – Transparency and Logic Involved.....	16
C.	Automated Decision-Making – Human Intervention and Right to Contest.....	17
V.	High-Level Expert Group Guidelines	17
VI.	Conclusion.....	19

Artificial Intelligence and Data Protection – How the GDPR Regulates AI

This paper forms a part of the Centre for Information Policy Leadership’s (CIPL)¹ special EU Project on Accountable AI. The project aims to facilitate expert dialogue and engagement between EU policy makers and businesses, leaders in AI use and development. The paper was written in collaboration with Olivia Lee, Associate at Hunton Andrews Kurth LLP

I. INTRODUCTION

On 19 February 2020 the European Commission published a “White Paper on Artificial Intelligence: a European approach to excellence and trust”² (the White Paper). This followed Ursula von der Leyen’s announcement that as president of the European Commission she intended to put forward legislation for a coordinated European approach on the human and ethical implications of artificial intelligence (AI) within her first 100 days in office (which commenced on 1 December 2019). The White Paper is open for consultation until 19 May 2020.

The European Commission, in its 2018 Communication on Artificial Intelligence, defined AI as “systems that display intelligent behaviour by analysing their environment and taking actions—with some degree of autonomy—to achieve specific goals.”³ This is just one of a number of descriptions of this increasingly used technology, which includes computer systems that perform tasks involving visual perception, speech recognition, decision-making and translation between languages.⁴ AI capabilities are advancing rapidly. However, in order to harness AI’s full potential, the legal concerns often raised regarding its use of personal data (i.e. information relating to an identified or identifiable individual) and the potential bias and unpredictability in its output need to be confronted.

The General Data Protection Regulation (EU) 2016/679 (GDPR) does not refer to AI specifically, but rather regulates the processing of personal data regardless of the technology used. As a consequence, any technology that is designed to process personal data, including AI, is fully captured by the regime. This includes many of the requirements for trustworthy AI as highlighted by the High-Level Expert Group on AI (HLEG), a group of 52 experts appointed by the EU Commission to support the implementation of its AI strategy.⁵ The European Data Protection Board (EDPB) commented in a recent letter to a Member of the

¹ CIPL is a global data privacy and cybersecurity think tank in the law firm of Hunton Andrews Kurth LLP and is financially supported by the law firm and 90 member companies that are leaders in key sectors of the global economy. CIPL’s mission is to engage in thought leadership and develop best practices that ensure both effective privacy protections and the responsible use of personal information in the modern information age. CIPL’s work facilitates constructive engagement between business leaders, privacy and security professionals, regulators and policymakers around the world. For more information, please see CIPL’s [website](#). Nothing in this note should be construed as representing the views of any individual CIPL member company or of the law firm of Hunton Andrews Kurth. This note has been prepared for informational purposes only and should not be construed as legal advice.

² [White Paper on Artificial Intelligence: a European approach to excellence and trust](#).

³ [Communication: Artificial Intelligence for Europe](#).

⁴ English Oxford Living Dictionaries, “Artificial Intelligence”, available at https://www.lexico.com/definition/artificial_intelligence.

⁵ HLEG: “[Ethics Guidelines for Trustworthy AI](#).”

European Parliament: “the GDPR is built in a technologically neutral manner in order to be able to face any technological change or revolution.”⁶

In addition, there are also GDPR provisions that specifically allude to technologies or methods of processing that incorporate aspects of AI. These include the GDPR’s provisions on automated decision-making. Finally, there are provisions of the GDPR that specifically address some of the common issues and risks associated with AI, such as those relating to the data protection principle of data minimisation and the requirement that personal data be processed lawfully and fairly. Myriad guidance has been published to date on the topic of how AI works within the context of data protection legislation, including by data protection authorities (DPAs)⁷ and by the HLEG.⁸

Below we examine how the GDPR already regulates AI systems. In Section II we detail how the GDPR applies generally to AI systems in the same manner as any other processing of personal data. In Section III we highlight the GDPR provisions of particular relevance in the context of AI, and how these provisions govern and limit its use. In Section IV we discuss the provisions of the GDPR that specifically regulate AI. Finally, in Section V we look at how the principles for trustworthy AI outlined by the HLEG overlap with the GDPR requirements.

II. HOW GDPR APPLIES GENERALLY IN THE CONTEXT OF AI

As stated by the EDPB, “[a]ny processing of personal data through an algorithm falls within the scope of the GDPR”.⁹ Therefore, whenever an AI system uses personal data, all of the standard provisions of the GDPR may apply. Where personal data is processed by an AI system, this is carried out in two distinct phases—the algorithmic training phase and the use phase. During the former the AI’s algorithm is trained on a set of data, allowing it to create a model by identifying patterns and connections between different data points. In the latter phase, this model is applied to the particular use case that the AI was designed for, in order to provide a prediction or classification, assist a human decision or make a decision itself. Personal data is therefore a vital component for the full life cycle of an AI system.

It should be noted that not all AI systems process personal data. But even where AI systems are not designed to process personal data, instead relying on anonymised data,¹⁰ the line between personal data and non-personal data is increasingly becoming blurred. This may stem from a lack of robustness in anonymisation techniques and the risk of re-identification stemming from the correlations and inferences that can be pulled from aggregated data sets. When it comes to AI in particular, the unforeseen consequences of using a system designed to make connections and spot patterns not immediately visible to the human eye increases this risk of re-identification.

⁶ EDPB [Response](#) to the MEP Sophie in’t Veld’s letter on unfair algorithms.

⁷ CNIL: [“How Can Humans Keep the Upper Hand? The ethical matters raised by algorithms and artificial intelligence”](#); AP: [“Toezicht op AI & Algoritmes”](#); ICO: [“Big Data, AI, Machine Learning and Data Protection”](#); AEPD: [“Una aproximación para la adecuación al RGPD de tratamientos que incorporan Inteligencia Artificial”](#); Datatilsynet: [“Artificial Intelligence and Privacy”](#).

⁸ See note 5.

⁹ EDPB [Response](#) to the MEP Sophie in’t Veld’s letter on unfair algorithms.

¹⁰ It should be noted that as per recital 26 and Article 4(5) GDPR, pseudonymised data – which is data no longer attributable to an individual without the use of additional information - is still considered to be personal data, hence subject to the GDPR.

A. Legal Basis

For all processing of personal data using AI systems, controllers¹¹ need to rely on one of the six legal bases for processing set out under Article 6(1) of the GDPR. Most commonly controllers rely on consent, legitimate interests, legal obligation or contractual necessity.¹² An appropriate legal basis should be established during both the training phase and the use phase.

There are various requirements associated with these legal bases. For example, for consent to be valid under the GDPR, it must be specific, informed, freely given and unambiguous.¹³ With respect to the legitimate interest legal basis, controllers are required to balance the interests they (or third parties) are pursuing and the interests and rights of the individuals whose personal data is being processed. In the AI context, this may mean defining the objective of the AI's processing at the outset and making sure the original purpose of the processing is re-evaluated if the AI system provides an unexpected result, either so that the legitimate interests pursued can be identified or so that valid consent, as the case may be, can be collected from individuals.

The GDPR sets a general prohibition on processing special categories of data, such as data relating to health, race or sexual orientation, except in specific circumstances, such as where the individual has provided explicit consent.¹⁴ Any use of AI to process such data, which is likely to be undertaken wherever AI is used in, for example, the health sector or in relation to crime prevention and detection, needs to rely on one of the specific derogations set out under Article 9 to this general prohibition.

The processing of personal data of children also warrants particular care under the GDPR. For example, where an information society service is offered to a child (i.e. someone under the age of 16) and his/her consent is obtained, that consent must be provided or authorised by an adult with parental responsibility for that child (though Member States are permitted to lower this age to 13).¹⁵ Controllers need to be more cautious when developing or using AI systems designed to offer such services to children.

B. Data Protection Principles

Article 5(1) of the GDPR sets out data protection principles that must be complied with under the GDPR. Some of these are of particular relevance to AI systems (see Section III), but all processing also need to

¹¹ In the context of AI there is some complexity in identifying which party acts as the “controller” for the purposes of the GDPR. Depending on the nature of the AI system, the purposes for which it is used, the stage at which it is used and the level of control each party has, the controller may be the designer of the system, the developer that trains it, the entity selling it or the entity using it.

¹² There may also be instances where AI is deployed in the public interest to protect the vital interests of individuals. The ICO states in its [guidance](#) that it may be difficult to prove that the use of technology such as big data analytics is “strictly necessary” in a contractual context.

¹³ The ICO has [suggested](#), as a way of meeting this standard, that it is possible to approach consent in a more granular fashion than regarding it as a yes/no binary choice. When it comes to AI in particular, it will be necessary to take a more flexible approach, where individuals provide consent to different forms of processing at different stages of a system's use of their data, as opposed to being provided only with a binary choice at the outset.

¹⁴ Article 9 GDPR.

¹⁵ Article 8(1) GDPR.

comply with the principles of purpose limitation, accuracy, storage limitation, and integrity and confidentiality (security).

The **purpose limitation principle** requires that controllers using AI systems determine the purpose of the AI system's use at the outset of its training or deployment, and perform a re-assessment of this determination should the system's processing throw up unexpected results, since it requires that personal data only be collected for "specified, explicit and legitimate purposes" and not used in a way that is incompatible with the original purpose.¹⁶ In the same vein, the **storage limitation principle** requires that personal data be kept in identifiable form for no longer than is necessary for the purposes for which the data is processed.¹⁷

The **accuracy principle** requires that personal data is accurate and kept up-to-date.¹⁸ This principle is of particular importance for fully automated systems, where the output could have a significant impact on individuals with little human oversight. Feeding an AI system inaccurate data could lower the quality of the output, and this principle requires that AI users take a particularly vigilant approach to ensuring that the data set is not diluted by bad quality data.¹⁹ This emphasis on the accuracy of data is highlighted by the EDPB in its draft Guidelines on Privacy by Design and by Default,²⁰ where it provides the example of a bank using AI to determine which customers are granted a loan. In an instance like this, where an individual may be relying on the decision that the algorithm makes, inaccurate data resulting in an imprecise decision could have a significant impact on individuals.

Finally, the principle that personal data should be **processed securely**²¹ requires that those developing, deploying or using AI consider the particular security risk issues that such use may raise and mitigate those issues promptly. In addition to potential unauthorised access to personal data, a lack of proper security may lead to unauthorised third parties' accessing and tampering with the algorithm to change its logic and outcomes. This may have serious consequences for individuals where, for example, a decision is made regarding them by or with the help of this algorithm.

C. Accountability

Accountability requires those processing personal data to put in place comprehensive organisational policies and procedures to ensure that personal data is processed in compliance with the GDPR's requirements and to be able to demonstrate those policies and procedures.²² In the context of AI, controllers need to be accountable to both regulators and individuals, and need to take into account the

¹⁶ Article 5(1)(b) GDPR.

¹⁷ Article 5(1)(e) GDPR.

¹⁸ Article 5(1)(d) GDPR.

¹⁹ It has been commented by the CNIL in its [report](#): "the temptation for negligence in this regard must be taken seriously. Especially in some areas where the impact of poor quality data might not be immediately perceptible, such as human resources and recruitment".

²⁰ [Guidelines 4/2019 on Article 25 Data Protection by Design and by Default](#) (version for public consultation).

²¹ Article 5(1)(f) GDPR.

²² Article 24 GDPR. The EDPB also stresses that "While Art. 24 GDPR primarily concerns the rights to data protection and privacy, it may also involve other fundamental rights such as freedom of speech, freedom of thought, freedom of movement, prohibition of discrimination, and the rights to liberty, conscience and religion." EDPB [Response](#) to the MEP Sophie in't Veld's letter on unfair algorithms.

likelihood and severity of the consequences of the AI's use on individuals. They cannot simply deploy an AI system and then blame that system when its output harms individuals or results in non-compliance.

Controllers are subject to more onerous accountability obligations than processors. In the context of AI, parties that would typically act as processors, such as software developers or system designers, may be acting as co-controllers (where each party decides the purposes and means of their own processing of the personal data used by the AI) or as joint controllers (e.g. where the parties develop an AI system together).²³ In these instances, each party needs to comply with the controllership obligations under the GDPR in relation to the AI system.

An organisation's accountability program needs to comprise several elements, as set out by CIPL's accountability wheel in its *White Paper on Organisational Accountability - Past, Present and Future*.²⁴ In the context of AI, an organisation's privacy management program includes: leadership and oversight (ensuring oversight and buy-in from top-level management on the need to develop, deploy or use AI responsibly); risk assessment (assessing the impact of the AI system on individuals to mitigate potential risks); the creation and implementation of appropriate policies and procedures (creating operational, verifiable and auditable controls); transparency (providing understanding, explainability, traceability and information on the benefits of the AI system and rights of individuals); training and awareness (ensuring that the relevant employees understand their responsibilities with respect to the development or use of a particular AI system); control and monitoring (verifying and auditing practices to uncover potential non-compliant situations related to the use of AI and to address them); and response and enforcement (responding to data breaches, individual complaints or inquiries from regulators).

With regard to the requirements for leadership and oversight, implementation of policies and procedures and response and enforcement, the French DPA (the CNIL) has commented that: "the roll-out of an algorithmic system systematically must give rise to a clear attribution of the liabilities that should be assumed in its operation". The CNIL recommends identifying a specific team or authority within a company that is responsible for the use of AI systems wherever personal data is processed, so that this team can be reached easily by individuals.²⁵

D. Appointing a Data Protection Officer

Controllers and processors must designate a data protection officer (DPO) where one of the following criteria applies:²⁶ (i) processing is carried out by a public authority or body;²⁷ (ii) their core activities consist

²³ The EJC has taken a broad approach to this classification—for example a party may be a joint controller even where it does not have access to any personal data processed by the system. The parties also do not need to share equal responsibility for the processing to be considered joint controllers. Case C-40/17 *Fashion ID GmbH & Co. KG v Verbraucherzentrale NRW eV*, [2020] 1 C.M.L.R. 16.

²⁴ CIPL report: "[CIPL White Paper on Organisational Accountability - Past, Present and Future](#)."

²⁵ CNIL: "[How Can Humans Keep the Upper Hand? The ethical matters raised by algorithms and artificial intelligence](#)."

²⁶ Article 37 GDPR.

²⁷ For example processing in the context of a government department.

of processing activities involving regular and systematic monitoring of individuals on a large scale;²⁸ or (iii) their core activities consist of processing on a large scale of special categories of data.²⁹

In the context of an AI system, there is a higher likelihood that one of these triggers is present than with other forms of processing. For example AI is used in the context of medical diagnoses (which will include sensitive personal data), and by its nature AI requires processing of large volumes of data, particularly in its training phase, to function effectively. As further examples, credit card issuers may use AI to detect and prevent financial fraud in millions of transactions, and law enforcement agencies may use facial recognition to assist with their surveillance activities.

E. Article 30 Inventory

Controllers are also required under Article 30 of the GDPR to maintain a record of processing that includes the purposes of each processing activity involving personal data, as well as the period for which the data is retained.³⁰ Controllers therefore need to be able to fill out this inventory record with all relevant information, including the purposes of their AI use at the outset of its training and deployment. If the output of the AI provides material that is inconsistent with those original purposes or if any other information changes through the life cycle of the system, the controller is required to update this information.

Processors are also required to maintain an Article 30 inventory record (with more limited detail) which includes a record of the details of the controller on behalf of whom they are processing data, the categories of processing, details of transfers outside the EEA and the technical and organisational security measures they have implemented. In the event that they train an algorithm or develop an AI system on the instructions of a controller, they would have to ensure that the relevant information is stored in the inventory record.

F. Controller-Processor Contracts and Data Transfers

The GDPR requires that the relationship between a controller and a third party service provider or processor is governed by a contract that includes specific data protection provisions ensuring continued protection of personal data and compliance with the GDPR.³¹ In the context of AI, the relationship between controller and processor may vary, depending on the precise roles and responsibilities the parties have in relation to the training and deployment of the AI system. These must be properly reflected in a contract. For example, a processor may train an algorithm under the instruction of a controller. Through the contractual provisions controllers can ensure that these AI systems are designed or operated to process personal data only in accordance with their instructions and for the purposes that are agreed between the parties.

²⁸ For example the tracking of the locations of a large population through wearable devices or travel cards.

²⁹ For example the provision of background checking services.

³⁰ Controllers must also record their name and contact details, a description of the categories of individuals and personal data, categories of recipients of that personal data, transfers of that data to third countries and a description of the technical and organisational security measures put in place.

³¹ Article 28 GDPR. These include placing limitations on how the processor can treat the data, who they may share it with and how they are required to assist the controller with its own GDPR compliance.

All data transfers from EU controllers and processors to any recipient outside of the EEA must be framed by a GDPR-compliant transfer mechanism, such as Standard Contractual Clauses or certification to the Privacy Shield (for transfers to the US), to ensure that recipients provide the same level of data protection as that required under the GDPR.³² Even where EU data is only used to train an AI system, those operating the system outside of the EEA will be held to the same standard of protection with regard to that data as the disclosing controllers or processors who are directly subject to the GDPR.

G. Data Breaches

Whenever there is a breach involving personal data processed by an AI system, whether in the training or in the use phase, the controller must inform DPAs and individuals of the breach, if the relevant conditions are met by the circumstances of the incident.³³ This requires AI users to ensure that they have proper visibility into the AI's functioning, in order to be able to identify breaches when they occur and properly mitigate them.

H. Individual Rights

Individuals have certain rights in relation to their personal data under the GDPR, including rights to access,³⁴ rectify or update their data,³⁵ request its deletion,³⁶ or restrict or object to the processing in question.³⁷ Further, individuals have the right to receive the personal data they have provided to a controller in a structured, commonly used and machine-readable format.³⁸ In order to comply with these rights, an AI system must be designed and operated so as to allow for controllers to identify and, as the case may be, retrieve the information requested by individuals. If certain conditions are met, controllers also need to ensure that they are able to erase or remove the personal data from the AI system.³⁹

I. Data Protection by Design and by Default

Article 25 of the GDPR requires that controllers implement appropriate technical and organisational measures to ensure effective implementation of the data protection principles and meet the requirements of the GDPR, both prior to and during processing activities. Therefore, those designing AI systems need to make sure the privacy principles and obligations outlined above, i.e. the requirement to provide individuals with the opportunity to exercise their rights, to have a record of processing, to establish a legal basis for processing, to process the data securely, etc. are considered at every stage of

³² Article 46 GDPR.

³³ Articles 33 and 34 GDPR—with regard to DPA notification, this is required where the breach is not unlikely to result in a risk to the rights and freedoms of natural persons. The threshold for notification to individuals is higher - the breach must be likely to result in a high risk.

³⁴ Article 15 GDPR.

³⁵ Article 16 GDPR.

³⁶ Article 17 GDPR.

³⁷ Articles 18 and 21 GDPR.

³⁸ Article 20 GDPR.

³⁹ This may be required where an individual withdraws consent to processing, the personal data are no longer necessary for the purpose for which they were collected or otherwise processed, the individual objects to the processing and there is no overriding legitimate interest (or an objection is made under Article 21(2)), the processing is unlawful, the personal data needs to be erased in compliance with a legal obligation or the personal data was collected from children through an information society service.

the system’s design. AI systems must be designed with data protection considerations in mind rather than relegating these considerations to the final stages of the system’s creation or use—data privacy must be a primary focus from the outset. In its draft Guidelines on Data Protection by Design and by Default, the EDPB refers to the collection of AI training data from data sources with correct and up-to-date information as an example of how to comply with this requirement. In addition, the EDPB notes that the AI system alone should not be relied on for compliance with the data protection principles, and that the reliability of results from the AI is checked at regular intervals.⁴⁰

Article 25 also requires that the amount of personal data collected, the extent of its processing and the period of its storage be, by default, limited to only what is necessary. This limits AI developers and users in terms of what personal data they can collect and process, and is unlikely to allow for controllers to collect or retain data on a “nice to have” basis.

J. Extraterritorial Effect of the GDPR

Under Article 3(2) of the GDPR, where an entity outside of the EU processes personal data of individuals in the EU either through offering those individuals goods or services, or monitoring their behaviour, that controller or processor is subject to GDPR and needs to appoint a representative in the EU.⁴¹ This representative is intended to provide both EU regulators and individuals with a contact point through which they can seek information regarding the AI’s use. This may be the case for instance where an AI system is trained on personal data collected through the monitoring of individuals in the EU or used in the context of offering certain goods or services.

K. Data Protection Authorities, Sanctions and Enforcement

DPA’s have turned their attention more broadly to AI and are providing specific guidance on its responsible use, including in France, Norway and the UK (all cited in this paper). In addition, the Dutch DPA, De Autoriteit Persoonsgegevens (AP), included AI and algorithms as a key focus for its supervision and enforcement priorities from 2019 to 2023.⁴² The Spanish DPA, the Agencia Española de Protección de Datos (AEPD), also published in February of 2020 a guide for those looking to make use of AI (including developers) setting out the privacy and quality guarantees that should be applied.⁴³ The German Federal Data Ethics Committee (Datenethikkommission) has provided recommendations to the Federal Government with regard to its AI strategy.⁴⁴ Finally, the EDPB is integrating AI-related examples in its guidelines, such as those on Data Protection by Design and by Default,⁴⁵ and AI is featured in its Work Plan for the year 2019-2020.⁴⁶

DPA’s are also picking up on specific areas where AI presents some risk to individual rights and freedoms. For example, in 2019 the UK’s Information Commissioner’s Office (ICO) called on the government to

⁴⁰ [Guidelines 4/2019 on Article 25 Data Protection by Design and by Default](#) (version for public consultation).

⁴¹ Article 27 GDPR.

⁴² AP: “[AP legt focus in toezicht op datahandel, digitale overheid en AI.](#)”

⁴³ AEPD: “[Una aproximación para la adecuación al RGPD de tratamientos que incorporan Inteligencia Artificial.](#)”

⁴⁴ Datenethikkommission: “[Recommendations of the Data Ethics Commission for the Federal Government’s Strategy on Artificial Intelligence.](#)”

⁴⁵ [Guidelines 4/2019 on Article 25 Data Protection by Design and by Default](#) (version for public consultation).

⁴⁶ [EDPB Work Program 2019-2020.](#)

introduce a statutory and binding code of practice on the deployment of Live Facial Recognition technology.⁴⁷ The CNIL also publicly called in November 2019 for a democratic debate on the challenges raised by facial recognition technologies.⁴⁸

The GDPR provides regulatory authorities with considerable enforcement powers, including the power to fine non-compliant companies up to 4 per cent of annual worldwide turnover, or 20 million euros (whichever is higher). DPAs have already turned their enforcement powers towards processing involving aspects of AI. For example, in 2017 the ICO determined that the Royal Free NHS Foundation Trust had failed to comply with the Data Protection Act 1998 in its provision of the personal data of 1.6 million patients to Google DeepMind (an AI company) without adequately informing the individuals (though this incident took place while the Data Protection Act 1998 was still in force).⁴⁹

Additionally, consumer groups and European courts have been paying attention to the use of AI. For example the Court of the Hague ruled in February of 2020 that the System Risk Indication (SyRI), an AI tool used by the Dutch government to combat fraud in benefits, allowances and taxes, did not comply with Article 8 of the European Convention on Human Rights (i.e. the right to respect for private and family life) and was disproportionate to the aims it sought to achieve.⁵⁰ In February 2020 a French court determined that schools using facial recognition technology had done so without a GDPR-compliant legal basis and in a disproportionate manner.⁵¹

Individuals themselves also have the ability to seek redress under Article 79 of the GDPR, which permits them to seek an effective judicial remedy against those processing their personal data if they believe that their rights under the GDPR have been infringed. Under Article 82 individuals also have the right to seek compensation in court from controllers or processors if they suffer material or non-material damage as a result of infringement of the GDPR.

These potential administrative sanctions and civil remedies,⁵² in combination with the individual rights discussed above, provide adequate protection under the GDPR for any individual who wishes to contest the use of their personal data by an AI system.

This demonstrates that there is already an extensive and robust regulatory system in place to regulate the AI that involves the processing of personal data, both through DPA oversight⁵³ and the ability of individuals

⁴⁷ ICO: "[Live facial recognition technology – police forces need to slow down and justify its use.](#)"

⁴⁸ CNIL: "[Facial recognition: for a debate living up to the challenges.](#)"

⁴⁹ ICO: "[Royal Free - Google DeepMind trial failed to comply with data protection law.](#)"

⁵⁰ NJCM cs/De Staat der Nederlanden; case C-09-550982-HA ZA 18-388

<https://uitspraken.rechtspraak.nl/inziendocument?id=ECLI:NL:RBDHA:2020:865>.

⁵¹ https://www.laquadrature.net/wp-content/uploads/sites/8/2020/02/1090394890_1901249.pdf.

⁵² Criminal sanctions may also be applicable in some countries. For instance in France, potential criminal penalties for non-compliance with data protection can be as high as 5 years in prison and a 300,000 euro fine for individuals, as well as a 1,500,000 euro fine for legal persons. Criminal offences created by the Data Protection Act 2018 in the UK allow for unlimited fines.

⁵³ See also EDPB [Response](#) to the MEP Sophie in 't Veld's letter on unfair algorithms: "This enforcement can take many forms, including, but not limited to: actively informing the public regarding their rights; engaging with stakeholders; informing and guiding organisations; assessing prior consultations and; carrying out investigations, which may lead to enforcement actions."

to hold those using AI to account through the courts. Furthermore, it is clear that the DPAs and courts are already using regularly their powers to enforce the cases where AI has resulted in a human rights violation.

III. GDPR PROVISIONS OF PARTICULAR RELEVANCE TO AI

As discussed above, the GDPR regulates AI to the same extent as any other tool/technology used for processing personal data. However, there are provisions of the GDPR that are of particular relevance to AI systems. These include:

- the requirement for processing to be fair (Article 5(1)(a));
- the principle of data minimisation (Article 5(1)(c));
- data protection impact assessments (Article 35).

A. Fair Processing

The concept of fair processing as set out under Article 5(1)(a) covers a number of processing practices and overlaps with the requirement for transparency with regard to AI systems. It also implies an analysis of whether the processing will impact adversely and unjustifiably the individuals involved.

As highlighted in CIPL's report *on Delivering Sustainable AI Accountability in Practice – Hard Issues and Practical Solutions*,⁵⁴ defining fairness is an ongoing challenge as it can cover a wide range of meanings. It is a subjective and contextual concept that is influenced by several social, cultural and legal factors and which is magnified in the AI context. The EDPB provides a definition of fairness in its Guidelines on Data Protection by Design and by Default. It states: "Fairness is an overarching principle which requires that personal data shall not be processed in a way that is detrimental, discriminatory, unexpected or misleading to the data subject."⁵⁵

Fair processing requires that controllers consider the likely impact of their use of AI on individuals⁵⁶ and continuously reassess it.⁵⁷ In particular, fair processing requires that AI systems do not produce bias. The HLEG's Guidelines provide that: "Data sets used by AI systems (both for training and operation) may suffer from the inclusion of inadvertent historical bias,⁵⁸ incompleteness and bad governance models. The continuation of such biases could lead to unintended (in)direct prejudice and discrimination."

It has also been highlighted by the Dutch DPA that the question of whether or not the output of an AI is "fair" is inextricably linked to both the circumstances at hand and subjective views on justice. Therefore the ability to explain the decision becomes even more important. According to the Dutch DPA, "[a]

⁵⁴ [CIPL Second AI Report: Hard Issues and Practical Solutions](#).

⁵⁵ [Guidelines 4/2019 on Article 25 Data Protection by Design and by Default](#) (version for public consultation).

⁵⁶ [CIPL Second AI Report: Hard Issues and Practical Solutions](#). Section II.A of the Report discusses whether impacts on society should also be included in this assessment.

⁵⁷ For examples of technical and procedural tools and frameworks that organisations develop to ensure fairness, see [CIPL Second AI Report: Hard Issues and Practical Solutions](#).

⁵⁸ Bias in the output of AI systems has been demonstrated in numerous studies, such as in systems that incorrectly predict a higher likelihood of recidivism by African Americans than Caucasians. ProPublica: "[Machine Bias](#)."

controller must actively account for and justify why an algorithm is fair and the use of the chosen algorithm does not lead to inappropriate results.”⁵⁹ If an AI system is not sufficiently transparent, it may be impossible for those overseeing its use to identify bias in its reasoning and output. In addition, in order for processing to be fair, controllers must comply with their obligations relating to purpose limitation and individual rights. This is discussed in more detail in Section II above.

B. Data Minimisation

The data minimisation principle set out under Article 5(1)(c) of the GDPR requires that personal data be “adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.”

By definition, AI systems need substantial amounts of data to operate effectively, particularly during the training phase. For example, an AI system looking at heart attack risk factors will be provided with medical data on heart disease and heart attacks as well as more general information from various patients’ medical records and their lifestyle (e.g. smoking, drinking, diabetes history) during the training phase. During its deployment and use, the data of a specific patient will be analysed within a framework created by the AI based on all the data processed during the training phase.

As such, AI systems may not be able to perform without first being trained on a large data set. While this may be viewed as creating tension between the use of AI systems and data protection law,⁶⁰ since it is not always possible to predict what data elements may be relevant to the objective of the system, the principle in itself does not limit the processing of data by way of reference to a specific volume or set of data elements—it refers to what is “necessary” for the purposes of the processing. What personal data is considered “necessary” varies depending on the AI system and the objective for which it is used, but the governance of the GDPR in this area should prevent the perfect from being the enemy of the good for AI designers - the fact that the personal data must be limited does not mean that the AI system itself becomes useless, especially since not all AI systems need to provide a precise output. As commented by the Norwegian DPA (Datatilsynet),⁶¹ the level of accuracy that is required is to be a determining factor in the selection of data elements for inclusion. For systems that require 100 per cent accuracy, for example in the medical sphere, controllers need to input more data than for systems where a margin for error is acceptable. The principle of data minimisation is a consideration during this analysis. The Norwegian DPA suggests as a practice that controllers should set limits that are sufficient to achieve the purpose of the processing, rather than using all available data.

An example of where an AI system has been designed with this in mind is a tool developed by the Norwegian Tax Administration in order to filter tax returns for errors. Five hundred variables were tested, but only thirty were included in the final AI model, as they proved most relevant to the task at hand. This demonstrates that compliance with the data minimisation principle (1) may require extensive processing of data in order to determine what is adequate, relevant and necessary and (2) can potentially assist controllers in streamlining their AI systems.

⁵⁹ AP: [“AP legt focus in toezicht op datahandel, digitale overheid en AI.”](#)

⁶⁰ CIPL report: [“Artificial Intelligence and Data Protection in Tension.”](#)

⁶¹ Datatilsynet: [“Artificial Intelligence and Privacy.”](#)

C. Data Protection Impact Assessments (DPIAs)

DPIAs are designed to assess the impact of a processing activity on the protection of personal data, where the processing is likely to result in a high risk to the rights and freedoms of natural persons.⁶² The DPIA needs to contain a systematic description of the proposed processing, its purpose and the legitimate interest pursued (if applicable), as well as an assessment of its necessity and proportionality, its risks and the measures envisaged to address those risks. Where the results of a DPIA demonstrate a high risk to individuals that cannot be mitigated, the controller is required to consult with a relevant DPA before carrying out the proposed processing.

Where AI is proposed to be used, controllers may be required to consider, prior to its training and/or prior to deployment, the risks that it poses to the individuals concerned and whether it has adequately mitigated those risks.

The use of AI is more likely to trigger the requirement for a DPIA, based on criteria in Art 35 GDPR. The GDPR and the EDPB's Guidelines on DPIAs identify both "new technologies" and the type of automated decision-making that produce legal effects or similarly significantly affect persons as likely to result in a "high risk to the rights and freedoms of natural persons".

New technologies include where technology is used innovatively, and where technologies are combined to augment their effect. These activities provide opportunities for new forms of data collection and usage according to the EDPB⁶³ and the risks of these novel activities are unknown. These new technologies are therefore highlighted as triggers for the requirement to carry out a DPIA.

The EDPB's Guidelines highlight the following additional processing activities as triggers for a DPIA, some of which are particularly relevant to AI:

- Evaluation or scoring (AI may be used to predict the likelihood of certain behaviours or outcomes, such as the likelihood of an applicant defaulting on a loan);
- Processing of personal data on a large scale (AI requires substantial input of data during its training and use phases); and
- Processing that prevents individuals from exercising a right or using a service or contract (AI may be deployed specifically to assist in making a decision regarding whether or not to enter into a contract with or provide a service to an individual).

The ICO also highlights AI, machine learning and deep learning as innovative technologies that likely trigger the requirement for a DPIA.⁶⁴ Similarly the CNIL has set out in its own DPIA Guidelines a number of processing activities relying on AI that may act as triggers for a DPIA, such as profiling for human resources purposes or profiling based on data collected from external sources.⁶⁵

⁶² Article 35 GDPR.

⁶³ [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, WP248 rev.01.](#)

⁶⁴ ICO: "[When do we need to do a DPIA?](#)"

⁶⁵ CNIL: "<https://www.cnil.fr/fr/liste-traitements-aipd-requise>"

Where consultation with a DPA is required because there is a residual risk in the system's use, the DPA has the final say on whether or not the AI's use is permissible. This may result in restriction of the entire AI system, or potentially only the aspects of the algorithm that are viewed as non-compliant.⁶⁶

IV. GDPR PROVISIONS THAT SPECIFICALLY REGULATE AI

The GDPR's provisions relating to automated decision-making (ADM), including profiling, specifically regulate AI.⁶⁷ As noted by the EDPB in its Guidelines on ADM:⁶⁸ "Advances in technology and the capabilities of big data analytics, artificial intelligence and machine learning have made it easier to create profiles and make automated decisions with the potential to significantly impact individuals' rights and freedoms." However, it is important to note that not all AI systems will automatically trigger the application of GDPR provisions on ADM.

A. Automated Decision-Making – Definition and Scope

In line with its risk-based approach, Article 22 of the GDPR subjects to additional requirements only those decisions based on automated processing and/or profiling that produce legal effect or a similarly significant effect for individuals. In addition, these decisions must be made solely by automated means, without a human intervention.

These additional protections therefore only apply to the most impactful ADM. The notion of "legal" effect and a "similarly significant" effect has been further analysed in *CIPL's comments on the EDPB's Guidelines on Automated Individual Decision-Making and Profiling*.⁶⁹ It requires a high threshold to cover only ADM producing an impact on someone's legal rights or something that affects a person's legal status or rights under a contract, or a decision with similar effects and significance. This covers for example ADM affecting accrued legal entitlements of a person or public rights, ADM affecting an individual's eligibility and access to essential services, admission to a country to a university or ADM to apply tax deductions.

Because Article 22 requires a high impact to apply, not all ADM will trigger additional protection. For instance ADM to ensure network security and prevent cyber-attacks, to reject fraudulent transactions in the context of fraud prevention, to disconnect a service when customers fail to make timely payments or commonly accepted forms of targeted advertising do not amount to producing a legal effect or a similarly significant effect for individuals and fall outside of the scope of Article 22.

Finally, Article 22 should not be considered in isolation as the only protection against ADM, but should be considered in combination with the other requirements and safeguards of the GDPR designed to protect individuals and ensure responsible use of data (see Sections II and III).

⁶⁶ EDPB [Response](#) to the MEP Sophie in't Veld's letter on unfair algorithms.

⁶⁷ Under Article 4(4) of the GDPR: "profiling means 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."

⁶⁸ [Guidelines on Automated Individual Decision-Making and Profiling](#).

⁶⁹ CIPL [comments](#) on the EDPB's [Guidelines on Automated Individual Decision-Making and Profiling](#).

B. Automated Decision-Making – Transparency and Logic Involved

Under Article 5(1)(a) of the GDPR, controllers are required to process personal data lawfully, fairly and **transparently**. Articles 13 and 14 also set out various notice requirements specifying what individuals should be informed of before their personal data is processed. In the context of AI use, these notice requirements include an obligation to inform individuals of the purposes for processing, their rights in relation to their data and the existence of ADM as discussed above, including meaningful information about the logic involved and the significance and envisaged consequences of such processing.⁷⁰

Similarly, Article 15 of the GDPR provides individuals with a right to access their data, which includes an obligation for the controller to provide information on the existence of ADM, including profiling, as well as meaningful information about the logic involved and the significance and envisaged consequences of the processing for the individual. Whether provided in a privacy notice or at the time of exercising the right of access, this information is of particular importance where AI is in use. Where automated processing and decision-making directly affect individuals, they will want to understand how the algorithm will come or has come to a decision concerning them.

The HLEG principles also highlight transparency, human agency, and oversight and accountability as three key principles for trustworthy AI, commenting that where a system may have a significant impact on people’s lives, “it should be possible to demand a suitable explanation of the AI system’s decision-making process.”⁷¹ The CNIL also includes in its six practical policy recommendations relating to AI that controllers make AI systems understandable by organising mediation with users.⁷²

Under GDPR, an AI system cannot be operated without a degree of oversight from the individuals whose personal data it is using. Organisations using AI must therefore be ready to coordinate with individuals on various aspects of the AI’s use, including at the outset by providing information regarding the logic used by the relevant algorithm and at a later stage in the response to an exercise of rights. However, transparency is a challenge in the context of AI as the information provided to the individual must remain simple and basic, in order to be meaningful, and adapted to the context. Algorithms are by definition complex and evolve over time. In addition, transparency is not unlimited as it should not lead to disclosing trade secrets or “helping individuals game the system.”⁷³

⁷⁰ A practical demonstration of how to comply with these requirements is shown in [Google’s Cloud AI Explanations](#), which makes use of counterfactuals. This involves the AI system auditing its own reasoning and analysis in order to ensure that its conclusions and predictions are properly supported, by quantifying the contribution of each data factor it has analysed to its final conclusion. These make the rationale employed by the AI system more easily understandable to those overseeing its use, and therefore more explainable to individuals, thus demonstrating compliance with the GDPR’s transparency requirements.

⁷¹ HLEG: “[Ethics Guidelines for Trustworthy AI.](#)”

⁷² CNIL: “[How Can Humans Keep the Upper Hand? The ethical matters raised by algorithms and artificial intelligence.](#)”

⁷³ [CIPL Second AI Report: Hard Issues and Practical Solutions.](#)

C. Automated Decision-Making – Human Intervention and Right to Contest

Although there is a general right to object to processing of personal data in certain circumstances under Article 21 of the GDPR, Article 22 sets out a more specific right for individuals not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly affects him or her (as defined in paragraph A).

Article 22(2) provides that solely ADM is permitted where the individual has explicitly consented,⁷⁴ where the processing is necessary to enter into a contract with the individual or where it is authorised by Member State law. However, these conditions do not allow for automated decision-making where special categories of data are involved (except in specified circumstances).⁷⁵

Article 22(3) also provides individuals with a right to obtain human intervention in a decision made by AI and the right to contest the decision. These are important rights for individuals in the context of AI that also ensure that the use of AI is fair and human-centric.

Allowing individuals to obtain transparent information, including an explanation of how a decision was made, is therefore key to enabling them to contest an automated decision under Article 22(3) with the controller, or the DPA, or in court.

V. HIGH-LEVEL EXPERT GROUP GUIDELINES

The HLEG’s Ethics Guidelines for Trustworthy AI⁷⁶ set out a framework for achieving Trustworthy AI, comprising the following seven key requirements that businesses should meet in designing AI systems:

- Human agency and oversight;
- Technical robustness and safety;
- Privacy and data governance;
- Transparency;
- Diversity, non-discrimination and fairness;
- Environmental and societal well-being; and
- Accountability.

These principles are not limited to data privacy and aim to address a broader set of concerns arising from AI. Yet, they overlap in various ways with the GDPR requirements and even draw from the GDPR approach

⁷⁴ For a different interpretation of Article 22(2), see CIPL [comments](#) on the EDPB’s [Guidelines on Automated Individual Decision-Making and Profiling](#).

⁷⁵ Article 22(4) GDPR.

⁷⁶ HLEG: “[Ethics Guidelines for Trustworthy AI](#).”

and concepts. For example they emphasise respect for human autonomy, and specifically human agency and associated rights as key requirements when using AI. According to the HLEG, “[u]sers should ... be given the knowledge and tools to comprehend and interact with AI systems to a satisfactory degree and, where possible, be enabled to reasonably self-assess or challenge the system ... The overall principle of user autonomy must be central to the system’s functionality.”⁷⁷

The table below shows the overlap between the HLEG Guidelines and the requirements of the GDPR, and the extent to which GDPR concepts have inspired the principles of trustworthy AI:

Key requirements of Trustworthy AI	Overlap with GDPR provisions
Human Agency and Oversight	<ul style="list-style-type: none"> • Legitimate interest balancing test (Art. 6(1)(f)) • Transparency (Art. 13 & 14) • ADM (Art. 22) and Right to obtain human intervention (Art. 22(3)) • Risk assessment and DPIA (Art. 35)
Technical Robustness and Safety	<ul style="list-style-type: none"> • Security (Art. 32) • Risk assessment and DPIA (Art. 35) • Data accuracy (Art. 5(1)(d))
Privacy and Data Governance	<ul style="list-style-type: none"> • Data protection principles (Art. 5) • Legal grounds for processing (Art. 6) • Legal grounds for sensitive data (Art. 9) • Rights of the data subject (Chapter III) and in particular Transparency (Art. 13 & 14); Right to information on ADM and logic involved (Art. 15(1)(h)); Right not to be subject to an ADM decision (Art. 22) and Right to human intervention (Art. 22(3)) • Accountability (Art. 5(2) & Art. 24(3)) • Data protection by design (Art. 25) • Processor due diligence (Art. 28(1)) • Security (Art. 32) • DPO (Art. 37 & 38)
Transparency	<ul style="list-style-type: none"> • Transparency (Art. 13 & 14) • ADM (Art. 22)

⁷⁷ HLEG: [“Ethics Guidelines for Trustworthy AI.”](#)

Key requirements of Trustworthy AI	Overlap with GDPR provisions
Diversity, Non-Discrimination and Fairness	<ul style="list-style-type: none"> • Fairness data protection principle (Art. 5.1(a)) • Risk assessment and DPIA (Art. 35) • Right to information on ADM and logic involved (Art. 15(1)(h))
Societal and Environmental Wellbeing	<ul style="list-style-type: none"> • Risk assessment and DPIA (Art. 35) • Transparency (Art. 13 & 14)
Accountability	<ul style="list-style-type: none"> • Accountability (Art. 5(2) & 24(3)) • Risk assessment and DPIA (Art. 35) • Processor due diligence (Art. 28(1)) • DPO (Art. 37 & 38)

VI. CONCLUSION

As demonstrated in this paper, the GDPR already extensively regulates AI. In the words of the EDPB: “[t]hanks to—inter alia—the risk based approach, the data minimisation principle and the requirement of data protection by design and by default, the current legal framework addresses many of the potential risks and challenges associated with the processing of personal data through algorithms”.⁷⁸ Thus, any regulatory regime that is designed to govern AI needs to take into consideration the GDPR, as well as its practical implementation by organisations, to avoid unnecessary duplication of existing and potentially conflicting obligations, ambiguity and legal uncertainty (particularly in heavily regulated sectors). In addition, to the extent any additional regulation specific to AI is contemplated or deemed necessary, it will need to be designed flexibly with an eye to the future, to anticipate the rapid progress and likely changes in this area. Any new AI law that is created without these important considerations in mind may do more harm for the development of and leadership on responsible AI than good.

⁷⁸ EDPB [Response](#) to the MEP Sophie in’t Veld’s letter on unfair algorithms.