

Opinion of the Board (Art. 64)



Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models

Adopted on 17 December 2024

Executive summary

AI technologies create many opportunities and benefits across a wide range of sectors and social activities.

By protecting the fundamental right to data protection, GDPR supports these opportunities and promotes other EU fundamental rights, including the right to freedom of thought, expression and information, the right to education or the freedom to conduct a business. In this way, GDPR is a legal framework that encourages responsible innovation.

In this context, taking into account the data protection questions raised by these technologies, the Irish supervisory authority requested the EDPB to issue an opinion on matters of general application pursuant to Article 64(2) GDPR. The request relates to the processing of personal data in the context of the development and deployment phases of Artificial Intelligence (“AI”) models. In more details, the request asked: (1) when and how an AI model can be considered as ‘anonymous’; (2) how controllers can demonstrate the appropriateness of legitimate interest as a legal basis in the development and (3) deployment phases; and (4) what are the consequences of the unlawful processing of personal data in the development phase of an AI model on the subsequent processing or operation of the AI model.

With respect to the first question, the Opinion mentions that claims of an AI model’s anonymity should be assessed by competent SAs on a case-by-case basis, since the EDPB considers that AI models trained with personal data cannot, in all cases, be considered anonymous. For an AI model to be considered anonymous, both (1) the likelihood of direct (including probabilistic) extraction of personal data regarding individuals whose personal data were used to develop the model and (2) the likelihood of obtaining, intentionally or not, such personal data from queries, should be insignificant, taking into account *‘all the means reasonably likely to be used’* by the controller or another person.

To conduct their assessment, SAs should review the documentation provided by the controller to demonstrate the anonymity of the model. In that regard, the Opinion provides a non-prescriptive and non-exhaustive list of methods that may be used by controllers in their demonstration of anonymity, and thus be considered by SAs when assessing a controller’s claim of anonymity. This covers, for instance, the approaches taken by controllers, during the development phase, to prevent or limit the collection of personal data used for training, to reduce their identifiability, to prevent their extraction or to provide assurance regarding state of the art resistance to attacks.

With respect to the second and third questions, the Opinion provides general considerations for SAs to take into account when assessing whether controllers can rely on legitimate interest as an appropriate legal basis for processing conducted in the context of the development and the deployment of AI models.

The Opinion recalls that there is no hierarchy between the legal bases provided by the GDPR, and that it is for controllers to identify the appropriate legal basis for their processing activities. The Opinion then recalls the three-step test that should be conducted when assessing the use of legitimate interest as a legal basis, i.e. (1) identifying the legitimate interest pursued by the controller or a third party; (2) analysing the necessity of the processing for the purposes of the legitimate interest(s) pursued (also referred to as “necessity test”); and (3) assessing that the legitimate interest(s) is (are) not overridden by the interests or fundamental rights and freedoms of the data subjects (also referred to as “balancing test”).

With respect to the first step, the Opinion recalls that an interest may be regarded as legitimate if the following three cumulative criteria are met: the interest (1) is lawful; (2) is clearly and precisely articulated; and (3) is real and present (i.e. not speculative). Such interest may cover, for instance, in the development of an AI model - developing the service of a conversational agent to assist users, or in its deployment - improving threat detection in an information system.

With respect to the second step, the Opinion recalls that the assessment of necessity entails considering: (1) whether the processing activity will allow for the pursuit of the legitimate interest; and (2) whether there is no less intrusive way of pursuing this interest. When assessing whether the condition of necessity is met, SAs should pay particular attention to the amount of personal data processed and whether it is proportionate to pursue the legitimate interest at stake, also in light of the data minimisation principle.

With respect to the third step, the Opinion recalls that the balancing test should be conducted taking into account the specific circumstances of each case. It then provides an overview of the elements that SAs may take into account when evaluating whether the interest of a controller or a third party is overridden by the interests, fundamental rights and freedoms of data subjects.

As part of the third step, the Opinion highlights specific risks to fundamental rights that may emerge either in the development or the deployment phases of AI models. It also clarifies that the processing of personal data that takes place during the development and deployment phases of AI models may impact data subjects in different ways, which may be positive or negative. To assess such impact, SAs may consider the nature of the data processed by the models, the context of the processing and the possible further consequences of the processing.

The Opinion additionally highlights the role of data subjects' reasonable expectations in the balancing test. This can be important due to the complexity of the technologies used in AI models and the fact that it may be difficult for data subjects to understand the variety of their potential uses, as well as the different processing activities involved. In this regard, both the information provided to data subjects and the context of the processing may be among the elements to be considered to assess whether data subjects can reasonably expect their personal data to be processed. With regard to the context, this may include: whether or not the personal data was publicly available, the nature of the relationship between the data subject and the controller (and whether a link exists between the two), the nature of the service, the context in which the personal data was collected, the source from which the data was collected (i.e., the website or service where the personal data was collected and the privacy settings they offer), the potential further uses of the model, and whether data subjects are actually aware that their personal data is online at all.

The Opinion also recalls that, when the data subjects' interests, rights and freedoms seem to override the legitimate interest(s) being pursued by the controller or a third party, the controller may consider introducing mitigating measures to limit the impact of the processing on these data subjects. Mitigating measures should not be confused with the measures that the controller is legally required to adopt anyway to ensure compliance with the GDPR. In addition, the measures should be tailored to the circumstances of the case and the characteristics of the AI model, including its intended use. In this respect, the Opinion provides a non-exhaustive list of examples of mitigating measures in relation to the development phase (also with regard to web scraping) and the deployment phase. Mitigating measures may be subject to rapid evolution and should be tailored to the circumstances of the case. Therefore, it remains for the SAs to assess the appropriateness of the mitigating measures implemented on a case-by-case basis.

With respect to the fourth question, the Opinion generally recalls that SAs enjoy discretionary powers to assess the possible infringement(s) and choose appropriate, necessary, and proportionate measures, taking into account the circumstances of each individual case. The Opinion then considers three scenarios.

Under scenario 1, personal data is retained in the AI model (meaning that the model cannot be considered anonymous, as detailed in the first question) and is subsequently processed by the same controller (for instance in the context of the deployment of the model). The Opinion states that whether the development and deployment phases involve separate purposes (thus constituting separate processing activities) and the extent to which the lack of legal basis for the initial processing activity impacts the lawfulness of the subsequent processing, should be assessed on a case-by-case basis, depending on the context of the case.

Under scenario 2, personal data is retained in the model and is processed by another controller in the context of the deployment of the model. In this regard, the Opinion states that SAs should take into account whether the controller deploying the model conducted an appropriate assessment, as part of its accountability obligations to demonstrate compliance with Article 5(1)(a) and Article 6 GDPR, to ascertain that the AI model was not developed by unlawfully processing personal data. This assessment should take into account, for instance, the source of the personal data and whether the processing in the development phase was subject to the finding of an infringement, particularly if it was determined by a SA or a court, and should be less or more detailed depending on the risks raised by the processing in the deployment phase.

Under scenario 3, a controller unlawfully processes personal data to develop the AI model, then ensures that it is anonymised, before the same or another controller initiates another processing of personal data in the context of the deployment. In this regard, the Opinion states that if it can be demonstrated that the subsequent operation of the AI model does not entail the processing of personal data, the EDPB considers that the GDPR would not apply. Hence, the unlawfulness of the initial processing should not impact the subsequent operation of the model. Further, the EDPB considers that, when controllers subsequently process personal data collected during the deployment phase, after the model has been anonymised, the GDPR would apply in relation to these processing operations. In these cases, the Opinion considers that, as regards the GDPR, the lawfulness of the processing carried out in the deployment phase should not be impacted by the unlawfulness of the initial processing.

Table of contents

- 1 Introduction..... 6
 - 1.1 Summary of facts..... 6
 - 1.2 Admissibility of the Request for an Article 64(2) GDPR opinion 8
- 2 Scope and key notions..... 9
 - 2.1 Scope of the Opinion 9
 - 2.2 Key notions 11
 - 2.3 AI models in the context of the Opinion 11
- 3 On the merits of the request..... 12
 - 3.1 On the nature of AI models in relation to the definition of personal data 12
 - 3.2 On the circumstances under which AI models could be considered anonymous and the related demonstration 14
 - 3.2.1 General consideration regarding anonymisation in the context at hand 14
 - 3.2.2 Elements to evaluate the residual likelihood of identification 16
 - 3.3 On the appropriateness of legitimate interest as a legal basis for processing of personal data in the context of the development and deployment of AI Models 19
 - 3.3.1 General observations 19
 - 3.3.2 Considerations on the three steps of the legitimate interest assessment in the context of the development and deployment of AI models 21
 - 3.4 On the possible impact of an unlawful processing in the development of an AI model on the lawfulness of the subsequent processing or operation of the AI model 31
 - 3.4.1 Scenario 1. A controller unlawfully processes personal data to develop the model, the personal data is retained in the model and is subsequently processed by the same controller (for instance in the context of the deployment of the model) 32
 - 3.4.2 Scenario 2. A controller unlawfully processes personal data to develop the model, the personal data is retained in the model and is processed by another controller in the context of the deployment of the model 33
 - 3.4.3 Scenario 3. A controller unlawfully processes personal data to develop the model, then ensures that the model is anonymised, before the same or another controller initiates another processing of personal data in the context of the deployment 34
- 4 Final remarks 35

The European Data Protection Board

Having regard to Article 63 and Article 64(2) of the Regulation 2016/679/EU 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 (hereinafter “GDPR”),

Having regard to the EEA Agreement and in particular to Annex XI and Protocol 37 thereof, as amended by the Decision of the EEA joint Committee No 154/2018 of 6 July 2018¹,

Having regard to Article 10 and Article 22 of its Rules of Procedure,

Whereas:

(1) The main role of the European Data Protection Board (hereafter the “**Board**” or the “**EDPB**”) is to ensure the consistent application of the GDPR throughout the European Economic Area (“**EEA**”). Article 64(2) GDPR provides that any supervisory authority (“**SA**”), the Chair of the Board or the Commission may request that any matter of general application or producing effects in more than one EEA Member State be examined by the Board with a view to obtaining an opinion. The aim of this opinion is to examine a matter of general application or which produces effects in more than one EEA Member State.

(2) The opinion of the Board shall be adopted pursuant to Article 64(3) GDPR in conjunction with Article 10(2) of the EDPB Rules of Procedure within eight weeks from when the Chair and the competent supervisory authority have decided that the file is complete. Upon decision of the Chair, this period may be extended by a further six weeks taking into account the complexity of the subject matter.

HAS ADOPTED THE FOLLOWING OPINION

1 Introduction

1.1 Summary of facts

1. On 4 September 2024, the Irish supervisory authority (the “**IE SA**” or “**requesting SA**”) requested the EDPB to issue an opinion pursuant to Article 64(2) GDPR in relation to AI models and the processing of personal data (“**the Request**”).
2. The Chair of the Board and the IE SA considered the file complete on 13 September 2024. On the following working day, 16 September 2024, the file was broadcast by the EDPB Secretariat. The Chair of the Board, considering the complexity of the matter, decided to extend the legal deadline, in line with Article 64(3) GDPR and Article 10(4) of the EDPB Rules of Procedure.
3. The Request addresses certain elements of the training, updating, developing and operation of AI models where personal data forms part of the relevant dataset. The IE SA highlights that the Request

¹ References to ‘Member States’ made throughout this opinion should be understood as references to ‘EEA Member States’. References to the ‘Union’ made throughout this opinion should be understood as references to the ‘EEA’.

concerns key issues that have a high impact on data subjects and controllers in the EEA, and that there is no harmonised position at this stage among the national SAs². The terminology that will be used for the purpose of this opinion is provided in Section 2.2 and 2.3 below.

4. The following questions were asked by the IE SA:

Question 1: Is the final AI Model, which has been trained using personal data, in all cases, considered not to meet the definition of personal data (as set out in Article 4(1) GDPR)?

If the answer to question 1 is “yes”:

- i. At what stage of the processing operations leading to an AI Model is personal data no longer processed?
 - a) How can it be demonstrated that the AI model does not process personal data?
- ii. Are there any factors which would cause the operation of the final AI Model to no longer be considered anonymous?
 - a) If so, how can the measures taken to mitigate, prevent or safeguard against these factors (so as to ensure the AI Model does not process personal data) be demonstrated?

If the answer to question 1 is “no”:

- i. What are the circumstances in which that might arise?
 - a) If so, how can the steps that have been taken to ensure that the AI Model is not processing personal data be demonstrated?

Question 2: Where a data controller is relying on legitimate interests as a legal basis for personal data processing to create, update and/or develop an AI Model, how should that controller demonstrate the appropriateness of legitimate interests as a legal basis, both in relation to the processing of third-party and first-party data?

- i. What considerations should that controller take into account to ensure that the interests of the data subjects, whose personal data are being processed, are appropriately balanced against the interests of that controller in the context of:
 - a) Third-party data
 - b) First-party data

Question 3: Post-training, where a data controller is relying on legitimate interests as a legal basis for personal data processing taking place within an AI Model, or an AI System of which an AI Model forms part, how should a controller demonstrate the appropriateness of legitimate interests as a legal basis?

Question 4: If an AI Model has been found to have been created, updated or developed using unlawfully processed personal data, what is the impact of this, if any, on the lawfulness of the continued or subsequent processing or operation of the AI model, either on its own or as part of an AI System, where:

² Request, p.1.

- i. The AI Model, either alone or as part of an AI System, is processing personal data?
- ii. Neither the AI Model, nor the AI Model as part of an AI System, is processing personal data?

1.2 Admissibility of the Request for an Article 64(2) GDPR opinion

5. Article 64(2) GDPR provides that, in particular, any SA may request that any matter of general application or producing effects in more than one Member State be examined by the Board with a view to obtaining an opinion.
6. The requesting SA addressed questions to the EDPB regarding data protection aspects in the context of AI models. It specified in the Request that, while many organisations are now using AI models, including large language models (“LLMs”), their operations, training and use raise ‘*a number of wide-ranging data protection concerns*’³, which ‘*impact data subjects across the EU/EEA*’⁴.
7. The Request raises, in essence, questions on (i) the application of the concept of personal data; (ii) the principle of lawfulness, with specific regard to the legal basis of legitimate interest, in the context of AI models; as well as, on (iii) the consequences of unlawful processing of personal data in the development phase of AI models, on the subsequent processing or operation of the model.
8. The Board considers that the Request concerns a ‘*matter of general application*’ within the meaning of Article 64(2) GDPR. In particular, the matter relates to the interpretation and application of Article 4(1), Article 5(1)(a) and Article 6 GDPR in relation to the processing of personal data in the development and deployment of AI models. As highlighted by the requesting SA, the application of these provisions to AI models raises systemic, abstract and novel issues⁵. The rapid development and deployment of AI models by more and more organisations raises specific issues and, as pointed out in the Request, ‘*the EDPB will greatly benefit from reaching a common position on the matters raised by this Request, such matters being central to the planned work of the EDPB in the short and medium term*’⁶. Additionally, AI technologies create many opportunities and benefits across a wide range of sectors and social activities. Besides, GDPR is a legal framework that encourages responsible innovation. It follows that a general interest exists in making this assessment in the form of an EDPB opinion, in order to ensure the consistent application of certain GDPR provisions in the context of AI models.
9. The alternative condition of Article 64(2) GDPR refers to matters ‘*producing effect in more than one Member State*’. The EDPB recalls that the term ‘*effects*’ is to be interpreted *lato sensu*, and hence is not simply limited to legal effects⁷. As more and more AI models are being trained and used by a growing number of organisations in the EEA, they do impact a large number of data subjects

³ Request, p.1.

⁴ Ibid.

⁵ Request, p. 2.

⁶ Request, p.1. As mentioned in the EDPB Work Programme for 2024-2025, adopted on 8 October 2024, available at https://www.edpb.europa.eu/system/files/2024-10/edpb_work_programme_2024-2025_en.pdf, the EDPB plans to issue, *inter alia*, guidelines on anonymisation, pseudonymisation, and data scraping in the context of generative AI.

⁷ EDPB, Internal document 3/2019 on Internal guidance on Article 64 (2) GDPR, adopted on 8 October 2019, paragraph 15, available at https://www.edpb.europa.eu/system/files/2022-07/internaledpb_document_201903_art64.2_en.pdf.

throughout the EEA, some of which have already raised concerns to their competent SA⁸. Therefore, the EDPB considers that the matter raised by the requesting SA also meets this condition.

10. The Request includes written reasoning on the background and motivations for submitting the questions to the Board, including on the relevant legal framework. Therefore, the Board considers that the Request is reasoned in line with Article 10(3) of the EDPB Rules of Procedure.
11. According to Article 64(3) GDPR⁹, the EDPB shall not issue an opinion if it has already issued one on the matter. The EDPB has not issued an opinion on the same matter and it has not yet provided replies to the questions arising from the Request.
12. For these reasons, the Board considers that the Request is admissible and the questions arising from it should be analysed in this opinion (the “**Opinion**”) adopted pursuant to Article 64(2) GDPR.

2 Scope and key notions

2.1 Scope of the Opinion

13. The Board agrees with the requesting SA that, from a data protection perspective, the development and deployment of AI models raise fundamental data protection questions. The questions relate in particular to: (i) when and how an AI Model can be considered as ‘anonymous’ (Question 1 of the Request); (ii) how can controllers demonstrate the appropriateness of legitimate interest as a legal basis in the development (Question 2 of the request) and deployment phases (Question 3 of the request); and (iii) whether the unlawful processing of personal data in the development phase has consequences on the lawfulness of the subsequent processing or operation of the AI model (Question 4 of the Request).
14. The EDPB recalls that SAs are responsible for monitoring the application of the GDPR and should contribute to its consistent application throughout the Union¹⁰. It is, therefore, within the competence of SAs to investigate specific AI models and, in doing so, to conduct case-by-case assessments.
15. This Opinion provides a framework for competent SAs to assess specific cases where (some of) the questions raised in the Request would arise. This Opinion does not aim to be exhaustive, but rather to provide general considerations on the interpretation of the relevant provisions, which competent SAs should take utmost account of when using their investigative powers. While this Opinion is addressed to competent SAs and relates to their activities and powers, it is without prejudice to the obligations of controllers and processors under the GDPR. In particular, pursuant to the accountability principle enshrined in Article 5(2) GDPR, controllers shall be responsible for, and be able to demonstrate compliance with, all the principles relating to their processing of personal data.
16. In some cases, some examples may be provided in the Opinion, but considering the broad scope of the questions included in the Request, as well as the different types of AI models covered therein, not all possible scenarios will be considered in this Opinion. Technologies associated with AI models are subject to rapid evolution; accordingly, the considerations of the EDPB in this Opinion should be interpreted in light of this.

⁸ Request, pp. 1-2.

⁹ Article 64(3) GDPR and Article 10(4) of the EDPB Rules of Procedure.

¹⁰ Article 51(1) GDPR and Article 51(2) GDPR.

17. **This Opinion does not analyse the below provisions, which may still play an important role when assessing the data protection requirements applicable to AI models:**

- **Processing of special categories of data:** The EDPB recalls the prohibition of Article 9(1) GDPR regarding the processing of special categories of data and the limited exceptions of Article 9(2) GDPR¹¹. In this respect, the Court of Justice of the European Union (“CJEU”) further clarified that *‘where a set of data containing both sensitive data and non-sensitive data is [...] collected en bloc without it being possible to separate the data items from each other at the time of collection, the processing of that set of data must be regarded as being prohibited, within the meaning of Article 9(1) of the GDPR, if it contains at least one sensitive data item and none of the derogations in Article 9(2) of that regulation applies’*¹². Furthermore, the CJEU also emphasised that *‘for the purposes of the application of the exception laid down in Article 9(2)(e) of the GDPR, it is important to ascertain whether the data subject had intended, explicitly and by a clear affirmative action, to make the personal data in question accessible to the general public’*¹³. These considerations should be taken into account when processing of personal data in the context of AI models involves special categories of data.
- **Automated-decision making, including profiling:** The processing operations conducted in the context of AI models may fall under the scope of Article 22 GDPR, which imposes additional obligations on controllers and provides additional safeguards to data subjects. The EDPB recalls, in this regard, its Guidelines on automated individual decision-making and profiling for the purposes of Regulation 2016/679¹⁴.
- **Compatibility of purposes:** Article 6(4) GDPR provides, for certain legal bases, criteria that a controller shall take into account to ascertain whether processing for another purpose is compatible with the purpose for which personal data are initially collected. This provision may be relevant in the context of the development and deployment of AI models and its applicability should be assessed by SAs.
- **Data protection impact assessments (“DPIAs”)** (Articles 35 GDPR): DPIAs are an important element of accountability, where the processing in the context of AI models is likely to result in a high risk to the rights and freedoms of natural persons¹⁵.
- **Principle of data protection by design** (Article 25(1) GDPR): Data protection by design is an essential safeguard to be assessed by SAs in the context of the development and deployment of an AI model.

¹¹ See also the EDPB Report of the work undertaken by the ChatGPT Taskforce, adopted on 23 May 2024, paragraph 18: ‘Regarding the processing of special categories of personal data, one of the exceptions of Article 9(2) must be applicable in addition, for the processing to be lawful. *In principle, one of these exceptions can be Article 9(2)(e) GDPR. However, the mere fact that personal data is publicly accessible does not imply that “the data subject has manifestly made such data public” [...]*’.

¹² CJEU judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 89.

¹³ CJEU judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 77.

¹⁴ Article 29 Working Party (“WP29”) Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, as last revised and adopted on 6 February 2018, endorsed by the EDPB on 25 May 2018. See also, CJEU judgment of 7 December 2023, Case C-634/21, *SCHUFA Holding and Others* (ECLI:EU:C:2023:957).

¹⁵ WP29 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, revised and adopted on 4 October 2017, endorsed by the EDPB on 25 May 2018.

2.2 Key notions

18. As a preliminary remark, the EDPB wishes to provide clarifications on the terminology and concepts it uses throughout this Opinion, and only for the purposes of this Opinion:
- “**First-party data**” refers to personal data which the controller has collected from the data subjects.
 - “**Third-party data**” refers to personal data that controllers have not obtained from the data subjects, but collected or received from a third party, for example from a data broker or collected via web scraping.
 - “**Web scraping**” is a commonly used technique for collecting information from publicly available online sources. Information scraped from, for example, services such as news outlets, social media, forum discussions and personal websites, may contain personal data.
 - The Request refers to the “**life-cycle of AI models**”, as well as to various stages regarding, inter alia, the ‘creation’, ‘development’, ‘training’, ‘update’, ‘fine-tuning’, ‘operation’ or ‘post-training’ of AI models. The EDPB acknowledges that, depending on the circumstances, such stages may take place in the development and deployment of AI models and may include the processing of personal data for various purposes of processing. Nevertheless, for the purpose of this Opinion, the EDPB considers important to streamline the categorisation of stages likely to occur. Therefore, for the purpose of this Opinion, the EDPB refers to the “**development phase**” and to the “**deployment phase**”. The development of an AI model covers all stages before any deployment of the AI model, and includes, inter alia, code development, collection of training personal data, pre-processing of training personal data, and training. The deployment of an AI model covers all stages relating to the use of an AI model and may include any operations conducted after the development phase. The EDPB remains aware of the variety of use-cases and of their potential consequences in terms of processing of personal data; thus, SAs should consider whether the observations provided in this Opinion are relevant for the processing they assess.
 - The EDPB also highlights that, when necessary, the term “**training**” refers to the part of the development phase where AI models learn from data to perform their intended task (as explained in the next Section of this Opinion).
 - The notion and scope of **AI models**, as it is understood by the EDPB for the purpose of this Opinion, is further specified in the following dedicated section.

2.3 AI models in the context of the Opinion

19. The EU Artificial Intelligence Act (“**AI Act**”)¹⁶ defines an ‘AI system’ as ‘*a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments*’¹⁷. Recital (12) of the AI Act further explains the notion of “AI system”. Accordingly, a key characteristic of AI systems is their capability to infer. The techniques that enable

¹⁶ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).

¹⁷ Article 3(1) AI Act.

inference while building an AI system include machine learning, logic- and knowledge-based approaches.

20. 'AI models', on the other hand, are only indirectly defined in the AI Act: *'Although AI models are essential components of AI systems, they do not constitute AI systems on their own. AI models require the addition of further components, such as for example a user interface, to become AI systems. AI models are typically integrated into and form part of AI systems'*¹⁸.
21. The EDPB understands that the definition of an AI model proposed in the Request is narrower than the one from the AI Act, as it refers to 'AI model' as *'to encompass the product resulting from the training mechanisms that are applied to a set of training data, in the context of Artificial Intelligence, Machine Learning, Deep Learning or other related processing contexts'* and further specifies that *'The term applies to AI Models which are intended to undergo further training, fine-tuning and/or development, as well as AI Models which are not.'*¹⁹
22. On that basis, the EDPB adopted this Opinion under the understanding that an AI system will rely on an AI model to perform its intended objective by incorporating the model into a larger framework (e.g. an AI system for customer service might use an AI model trained on historical conversation data to provide responses to user queries).
23. Furthermore, AI models (or **"models"**) relevant for this Opinion are those developed through a training process. Such training process is a part of the development phase, where models learn from data to perform their intended task. Therefore, the training process requires a dataset from which the model will identify and 'learn' patterns. In these cases, the model will use different techniques to build a representation of the knowledge extracted from the training dataset. This is notably the case with machine learning.
24. In practice, any AI model is an algorithm, whose functioning is determined by a set of elements. For example, deep learning models are often in the form of a neural network with multiple layers consisting of nodes connected by edges that have weights, which are adjusted during training to learn the relationships between inputs and outputs. The characteristics of a simple deep learning model would be: (i) the type and size of each layer, (ii) the weight attributed to each edge (sometimes called 'parameters'), (iii) the activation functions²⁰ between layers, and possibly (iv) other operations that may happen between layers. For instance, when training a simple deep learning model for image classification, inputs (the **"image pixels"**) will be associated with outputs, and weights may be adjusted, so as to produce the right output most of the time.
25. Other examples of deep learning models include LLMs and generative AI, which are used for e.g. generating human-like content and creating new data.
26. **Based on the above considerations, in line with the Request, the scope of this Opinion only covers the subset of AI models that are the result of a training of such models with personal data.**

3 On the merits of the request

3.1 On the nature of AI models in relation to the definition of personal data

¹⁸ Recital 97 AI Act.

¹⁹ Request, p. 3.

²⁰ I.e. functions that calculate, based on inputs and weights, the output of a neural node that will then be sent to the next layer of the neural network.

27. Article 4(1) GDPR defines personal data as *'any information relating to an identified or identifiable natural person'* (i.e., the data subject). Furthermore, Recital 26 GDPR provides that data protection principles should not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person, taking into account *'all the means reasonably likely to be used'* by the controller or another person. This includes: (i) data that was never related to an identified or identifiable individual; and (ii) personal data which has been rendered anonymous in such a manner that the data subject is not or no longer identifiable.
28. Accordingly, Question 1²¹ of the Request can be answered by analysing if an AI model resulting from training which involves processing of personal data should, in all cases, be considered anonymous. Based on the phrasing of the question, the EDPB will refer in this section to the process of 'training' an AI model.
29. First and foremost, the EDPB would like to provide the following general considerations. AI models, regardless of whether they are trained with personal data or not, are usually designed to make predictions or draw conclusions, i.e. they are designed to infer. Furthermore, AI models trained with personal data are often designed to make inferences about individuals different from those whose personal data were used to train the AI model. However, some AI models are specifically designed to provide personal data regarding individuals whose personal data were used to train the model, or in some way to make such data available. In these cases, such AI models will inherently (and typically necessarily) include information relating to an identified or identifiable natural person, and so will involve the processing of personal data. Therefore, these types of AI models cannot be considered anonymous. This would be the case, for example, (i) of a generative model fine-tuned on the voice recordings of an individual to mimic their voice; or (ii) any model designed to reply with personal data from the training when prompted for information regarding a specific person.
30. Based on the above considerations, in answering Question 1 of the Request, the EDPB focuses on the situation of AI models that are not designed to provide personal data related to the training data.
31. The EDPB considers that, even when an AI model has not been intentionally designed to produce information relating to an identified or identifiable natural person from the training data, information from the training dataset, including personal data, may still remain 'absorbed' in the parameters of the model, namely represented through mathematical objects. They may differ from the original training data points, but may still retain the original information of those data, which may ultimately be extractable or otherwise obtained, directly or indirectly, from the model. Whenever information relating to identified or identifiable individuals whose personal data was used to train the model may be obtained from an AI model with means reasonably likely to be used, it may be concluded that such a model is not anonymous.
32. In this regard, the Request states that *'Existing research publications highlight some potential vulnerabilities that can exist in AI Models which could result in personal data being processed,²² as well as the personal data processing that may go on when models are deployed for use with other data, either through Application Programming Interfaces ("APIs") or "prompt" interfaces²³*.

²¹ *'Is the final AI Model, which has been trained using personal data, in all cases, considered not to meet the definition of personal data (as set out in Article 4(1) GDPR)?'*

²² Such as Membership Inference Attacks ([OWASP](#)), and Model Inversion Attacks ([OWASP](#) & [Veale et al](#), 2018).

²³ Request, p.1-2.

33. In the same vein, research on training data extraction is particularly dynamic²⁴. It shows that it is possible, in some cases, to use means reasonably likely to extract personal data from some AI models, or simply to accidentally obtain personal data through interactions with an AI model (for instance as part of an AI system). Continuous research efforts in this field will help assessing further the residual risks of regurgitation²⁵ and extraction of personal data in any given case.
34. **Based on the above considerations, the EDPB considers that AI models trained on personal data cannot, in all cases, be considered anonymous. Instead, the determination of whether an AI model is anonymous should be assessed, based on specific criteria, on a case-by-case basis.**

3.2 On the circumstances under which AI models could be considered anonymous and the related demonstration

35. Regarding Question 1 of the Request²⁶, the EDPB is requested to clarify the circumstances in which an AI model, which has been trained using personal data, may be considered as anonymous. With regard to Question 1(i)(a) of the Request²⁷, the EDPB is requested to clarify which proof and/or documentation SAs should take into account when assessing whether an AI model is anonymous.

3.2.1 General consideration regarding anonymisation in the context at hand

36. The use of the expression ‘*any information*’ in the definition of ‘*personal data*’ within Article 4(1) GDPR reflects the aim to assign a wide scope to that concept, that encompasses all kinds of information provided that it ‘*relates*’ to the data subject, who are identified or can be identified directly or indirectly.
37. Information may relate to a natural person even when it is technically organised or encoded (for instance in an only machine-readable format, whether proprietary or open) in a way that does not make the relation with that natural person immediately apparent. In such cases, software applications may be used to easily identify, recognise and extract specific data. This is particularly true for AI models where parameters represent statistical relationships between the training data, and where it may be

²⁴ See, in this regard, for instance: (i) Veale M., Binns R., Edwards L., 2018, *Algorithms that remember: model inversion attacks and data protection law*. Phil. Trans. R. Soc. A 376: 20180083, available at <http://dx.doi.org/10.1098/rsta.2018.0083>; (ii) Brown H., Lee K., Mireshghallah F., Shokri R., and Tramèr F., *What Does it Mean for a Language Model to Preserve Privacy?*, 2022, ACM Digital Library, FAccT '22, June 20, 2022, Seoul, Republic of Korea, available at <https://dl.acm.org/doi/abs/10.1145/3531146.3534642>; (iii) Vassilev A., Oprea A., Fordyce A., Anderson H., *Adversarial Machine Learning A Taxonomy and Terminology of Attacks and Mitigations*, January 2024, National Institute of Standards and Technology, available at <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2023.pdf>; (iv) Carlini N., Tramèr F., Wallace E., Jagielski M., Herbert-Voss A., Lee K., Roberts A., Brown T., Song D., Erlingsson U., Oprea A., Raffel C., *Extracting Training Data from Large Language Models*, arXiv:2012.07805v2 [cs.CR] 15 Jun 2021, available at <https://arxiv.org/pdf/2012.07805>; (v) Fredrikson M., Jha S., Ristenpart T., *Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures*, ACM Digital Library, 12 October 2015, available at <https://dl.acm.org/doi/abs/10.1145/2810103.2813677>; (vi) Zhang Y., Jia R., Pei H., Wang W., Li B., Song D., *The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks*, arXiv:1911.07135v2 [cs.LG] 18 Apr 2020, available at <https://arxiv.org/pdf/1911.07135>.

²⁵ For an AI system based on generative AI, regurgitation corresponds to the situation where outputs would directly relate to training data.

²⁶ ‘*What are the circumstances in which that might arise?*’

²⁷ ‘*If so, how can the steps that have been taken to ensure that the AI Model is not processing personal data be demonstrated?*’

possible to extract accurate or inaccurate (because statistically inferred) personal data, either directly from the relationships between the data included in the model, or by querying that model.

38. As AI models usually do not contain records that may be directly isolated or linked, but rather parameters representing probabilistic relationships between the data contained in the model, it may be possible to infer²⁸ information from the model, such as membership inference, in realistic scenarios. Therefore, for a SA to agree with the controller that a given AI model may be considered anonymous, it should check at least whether it has received sufficient evidence that, with reasonable means: (i) personal data, related to the training data, cannot be extracted²⁹ out of the model; and (ii) any output produced when querying the model does not relate to the data subjects whose personal data was used to train the model.
39. Three elements should be considered by SAs in the assessment of whether these conditions are fulfilled.
40. First, SAs should consider the elements identified in the most recent WP29 opinions and/or EDPB guidelines on the matter. Regarding anonymisation at the date of this Opinion, SAs should consider the elements included in the WP29 Opinion 05/2014 on Anonymisation Techniques (the “**WP29 Opinion 05/2014**”), which states that if it is not possible to single out, link and infer information from the supposedly anonymous dataset, the data may be considered anonymous³⁰. It also states that, ‘*whenever a proposal does not meet one of the criteria, a thorough evaluation of the identification risks should be performed*’³¹. **Given the above-mentioned likelihood of extraction and inference, the EDPB**

²⁸ (i) Carlini N., Chien S., Nasr M., Song S., Terzis A., Tramer F., *Membership Inference Attacks From First Principles*, arXiv:2112.03570, available at <https://arxiv.org/abs/2112.03570>;

(ii) Crețu A.M, Guépin F., and De Montjoye Y.A, *Correlation inference attacks against machine learning models*. Sci. Adv.10, eadj9260(2024). DOI:10.1126/sciadv.adj9260 available at <https://www.science.org/doi/10.1126/sciadv.adj9260>;

(iii) Dana L., Pydi M. S., Chevalyre Y., *Memorization in Attention-only Transformers* arXiv:2411.10115v1 [cs.AI] 15 November 2024, available at: <https://arxiv.org/abs/2411.10115>;

(iv) Gehrke M., Liebenow J., Mohammadi E. & Braun T. et al. *Lifting in Support of Privacy-Preserving Probabilistic Inference*. *Künstl Intell*, 13 June 2024, available at: <https://doi.org/10.1007/s13218-024-00851-y>;

(v) Hu H., *Membership Inference Attacks and Defenses on Machine Learning Models Literature*, available at: <https://github.com/HongshengHu/membership-inference-machine-learning-literature>;

(vi) Nasr M., Carlini N., Hayase J., Jagielski M., Cooper A. F., Ippolito D., Choquette-Choo C. A., Wallace E., Tramèr F., and Lee K., *Scalable Extraction of Training Data from (Production) Language Models*, arXiv:2311.17035 28 November 2023, available at: <https://arxiv.org/abs/2311.17035>;

(vii) Shokri R., Stronati M., Song C., Shmatikov V., *Membership Inference Attacks against Machine Learning Models* arXiv:1610.05820v2 [cs.CR], 31 March 2017, available at <https://arxiv.org/abs/1610.05820>;

(viii) Staab R., Vero M., Mislav Balunović, Martin Vechev, 2024, *Beyond Memorization: Violating Privacy Via Inference with Large Language Models*, arXiv:2310.07298v2, 6 May 2024, available at <https://arxiv.org/abs/2310.07298>;

(ix) Wu F., Cui L., Yao S., Yu S., *Inference Attacks in Machine Learning as a Service: A Taxonomy, Review, and Promising Directions* arXiv:2406.02027v1 [cs.LG], 27 June 2024, available at <https://arxiv.org/abs/2406.02027v1>;

(x) Zhang J., Das D., Kamath G., Tramèr F., *Membership Inference Attacks Cannot Prove that a Model Was Trained On Your Data* arXiv:2409.19798v1, [cs.LG], 29 September 2024, available at <https://arxiv.org/abs/2409.19798>;

(xi) Zhou Z., Xiang J., Chen C., and Su S., *Quantifying and Analyzing Entity-Level Memorization in Large Language Models*, arXiv:2308.15727v2 [cs.CL] 5 Nov 2023, available at: <https://arxiv.org/abs/2308.15727>.

²⁹ Extraction includes in particular the case where personal data is deduced from the AI model itself, with little or no use of the query interfaces.

³⁰ WP29 Opinion 05/2014, p.24.

³¹ WP29 Opinion 05/2014, p.24.

considers that AI models are very likely to require such a thorough evaluation of the risks of identification.

41. Second, this assessment should be made taking into account *'all the means reasonably likely to be used'* by the controller or another person to identify individuals³², and the determination of those means should be based on objective factors, as explained in Recital 26 GDPR, which may include:
 - a. the characteristics of the training data itself, the AI model and the training procedure³³;
 - b. the context in which the AI model is released and/or processed³⁴;
 - c. the additional information that would allow the identification and may be available to the given person;
 - d. the costs and amount of time that the person would need to obtain such additional information (in case it is not already available to them)³⁵; and
 - e. the available technology at the time of the processing, as well as technological developments³⁶.
42. Third, SAs should consider whether controllers have assessed the risk of identification by the controller and by different types of *'other persons'*, including unintended third parties accessing the AI model, also considering whether they can reasonably be considered to be able to gain access or process the data in question.
43. **In sum, the EDPB considers that, for an AI model to be considered anonymous, using reasonable means, both (i) the likelihood of direct (including probabilistic) extraction of personal data regarding individuals whose personal data were used to train the model; as well as (ii) the likelihood of obtaining, intentionally or not, such personal data from queries, should be insignificant³⁷ for any data subject. By default, SAs should consider that AI models are likely to require a thorough evaluation of the likelihood of identification to reach a conclusion on their possible anonymous nature. This likelihood should be assessed taking into account *'all the means reasonably likely to be used'* by the controller or another person, and should also consider unintended (re)use or disclosure of the model.**

3.2.2 Elements to evaluate the residual likelihood of identification

44. While measures might be taken both at the development and the deployment stages in order to reduce the likelihood of obtaining personal data from an AI model, the evaluation of anonymity of an AI model should also consider direct access to the model.
45. In addition, SAs should evaluate, on a case-by-case basis, if the measures implemented by the controller to ensure and prove that an AI model is anonymous are appropriate and effective.

³² CJEU judgment of 19 October 2016, Case C-582/14, *Breyer v Bundesrepublik Deutschland* (ECLI:EU:C:2016:779), paragraph 43.

³³ This includes characteristics such as uniqueness of the records in the training data, precision of the information, aggregation, randomization, and in particular how these affect the vulnerability to identification.

³⁴ This includes contextual elements, such as limiting access only to some persons and legal safeguards.

³⁵ CJEU judgment of 7 March 2024, Case C-479/22 P, *OC v European Commission* (ECLI:EU:C:2024:215), paragraph 50.

³⁶ CJEU judgment of 7 March 2024, Case C-479/22 P, *OC v European Commission* (ECLI:EU:C:2024:215), paragraph 50.

³⁷ CJEU judgment of 19 October 2016, Case C-582/14, *Breyer v Bundesrepublik Deutschland* (ECLI:EU:C:2016:779), paragraph 46, and CJEU judgment of 7 March 2024, Case C-479/22 P, *OC v European Commission* (ECLI:EU:C:2024:215), paragraph 51.

46. In particular, the conclusion of a SA's assessment might differ between a publicly available AI model, which is accessible to an unknown number of people with an unknown range of methods to try and extract personal data, and an internal AI model only accessible to employees. While in both cases SAs should verify that controllers have fulfilled their accountability obligation under Article 5(2) and Article 24 GDPR, the *'means reasonably likely to be used'* by other persons may have an impact on the range and nature of the possible scenarios to be considered. Therefore, depending on the context of development and deployment of the model, SAs may consider different levels of testing and resistance to attacks.
47. In that regard, the EDPB provides below a non-prescriptive and non-exhaustive list of possible elements that may be considered by SAs when assessing a controller's claim of anonymity. Other approaches may be possible if they offer an equivalent level of protection, in particular taking into account the state of the art.
48. The presence or absence of the elements listed below is not a conclusive criterion for assessing the anonymity of an AI model.

3.2.2.1 AI model design

49. Regarding AI model design, SAs should evaluate the approaches taken by controllers during the development phase. The application and effectiveness of four key areas (identified below) should be considered in this regard.

Selection of sources

50. The first evaluation area involves examining the selection of sources used to train the AI model. This includes an evaluation, by SAs, of any steps taken to avoid or limit the collection of personal data, including, among other things, (i) the appropriateness of the selection criteria; (ii) the relevance and adequacy of the chosen sources considering the intended purpose(s); and (iii) whether inappropriate sources have been excluded.

Data Preparation and Minimisation

51. The second area of evaluation relates to the preparation of data for the training phase. SAs should examine in particular: (i) whether the use of anonymous and/or personal data that has undergone pseudonymisation have been considered; and (ii) where it was decided not to use such measures, the reasons for this decision, taking into account the intended purpose; (iii) the data minimisation strategies and techniques employed to restrict the volume of personal data included in the training process; and (iv) any data filtering processes implemented prior to model training intended to remove irrelevant personal data.

Methodological choices regarding the training

52. The third area of evaluation concerns the selection of robust methods in AI model development. SAs should assess methodological choices that may significantly reduce or eliminate the identifiability, including, among others: (i) whether that methodology uses regularisation methods to improve model generalisation and reduce overfitting; and, crucially, (ii) whether the controller implemented appropriate and effective privacy-preserving techniques (e.g. differential privacy).

Measures regarding outputs of the model

53. The last area of evaluation concerns any methods or measures added to the AI model itself that may not have an impact on the risk of direct extraction of personal data for the model by anyone accessing it directly, but which might lower the likelihood of obtaining personal data related to training data from queries.

3.2.2.2 AI model analysis

54. For SAs to assess the robustness of the designed AI model regarding anonymisation, a first step is to ensure that the design has been developed as planned and is subject to effective engineering governance. SAs should evaluate whether controllers have conducted any document-based audits (internal or external) that include an evaluation of the chosen measures and of their impact to limit the likelihood of identification. This could include the analysis of reports of code reviews, as well as a theoretical analysis documenting the appropriateness of the measures chosen to reduce the likelihood of re-identification of the concerned model.

3.2.2.3 AI model testing and resistance to attacks

55. Finally, SAs should take into consideration the scope, frequency, quantity and quality of tests that the controller has conducted on the model. In particular, SAs should take into account that successful testing which covers widely known, state-of-the-art attacks can only be evidence for the resistance to those attacks. At to the date of this Opinion, this could include, among others, structured testing against: (i) attribute and membership inference; (ii) exfiltration; (iii) regurgitation of training data; (iv) model inversion; or (v) reconstruction attacks.

3.2.2.4 Documentation

56. Articles 5, 24, 25, and 30 GDPR, and, in cases of likely high risk to the rights and freedoms of data subjects, Article 35 GDPR, require controllers to adequately document their processing operations. This also applies to any processing that would include the training of an AI model, even if the objective of the processing is anonymisation. SAs should consider such documentation and any regular assessment of the consequential risks for the processing carried out by controllers, as they are fundamental steps to demonstrate that personal data is not processed.
57. **The EDPB considers that SAs should take into account the documentation whenever a claim of anonymity regarding a given AI model needs to be evaluated. The EDPB notes that, if a SA is not able to confirm, after assessing the claim of anonymity, including in light of the documentation, that effective measures were taken to anonymise the AI model, the SA would be in a position to consider that the controller has failed to meet its accountability obligations under Article 5(2) GDPR. Therefore, compliance with other GDPR provisions should also be considered.**
58. Ideally, SAs should verify whether the controller's documentation includes:
- a. any information relating to DPIAs, including any assessments and decisions that determined that a DPIA was not necessary;
 - b. any advice or feedback provided by the Data Protection Officer ("DPO") (where a DPO was - or should have been - appointed);
 - c. information on the technical and organisational measures taken while designing the AI model to reduce the likelihood of identification, including the threat model and risk assessments on which these measures are based. This should include the specific measures for each source of training datasets, including relevant source URLs and descriptions of measures taken (or already taken by third-party dataset providers);
 - d. the technical and organisational measures taken at all stages throughout the lifecycle of the model, which either contributed to, or verified the lack of personal data in the model;
 - e. the documentation demonstrating the AI model's theoretical resistance to re-identification techniques, as well as the controls designed to limit or assess the success and impact of main attacks (regurgitation, membership inference attacks, exfiltration, etc.). This may include, in

particular: (i) the ratio between the amount of training data, and the number of parameters in the model, including the analysis of its impact on the model³⁸; (ii) metrics on the likelihood of re-identification based on the current state-of-the-art; (iii) reports on how the model has been tested (by whom, when, how and to which extent) and (iv) the results of the tests;

- f. the documentation provided to the controller(s) deploying the model and/or to data subjects, in particular the documentation related to the measures taken to reduce the likelihood of identification and regarding the possible residual risks.

3.3 On the appropriateness of legitimate interest as a legal basis for processing of personal data in the context of the development and deployment of AI Models

59. To answer Questions 2 and 3 of the Request, the EDPB will first provide general observations on some important aspects that SAs should take into account, regardless of the legal basis for processing, when assessing how controllers may demonstrate compliance with the GDPR in the context of AI models. The EDPB, building on the Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR³⁹, will then consider the three steps required by the legitimate interest assessment in the context of the development and deployment of AI models.

3.3.1 General observations

60. The EDPB recalls that the GDPR does not establish any hierarchy between the different legal bases laid down in Article 6(1) GDPR⁴⁰.
61. Article 5 GDPR sets the principles relating to the processing of personal data. The EDPB highlights those that are significant for this Opinion and should at least be considered by SAs when assessing specific AI models, as well as the most relevant requirements from other provisions of the GDPR, taking into consideration the scope of this Opinion.
62. **Accountability principle** (Article 5(2) GDPR) - This principle provides that the controller shall be responsible for, and be able to demonstrate, compliance with the GDPR. In this regard, the roles and responsibilities of the parties that process personal data in the context of the development or deployment of an AI model should be assessed before the processing takes place, in order to define the obligations of the controllers or joint controllers, and of processors (if any), from the outset.
63. **Lawfulness, fairness and transparency principles** (Article 5(1)(a) GDPR) - When assessing the lawfulness of the processing in the context of AI models, in light of Article 6(1) GDPR, the EDPB considers it useful to distinguish the different stages of the processing of personal data⁴¹. The principle of fairness, which is closely related to the principle of transparency, requires that personal data is not processed by unfair methods, or by deception, or in a way that is '*unjustifiably detrimental, unlawfully*

³⁸ Ricciato F., *A Cautionary Reflection on (Pseudo-)Synthetic Data from Deep Learning on Personal Data*, Privacy in Statistical Databases conference (PSD 2024), Antibes, France, September 2024, slides available at: https://cros.ec.europa.eu/system/files/2024-10/20240926_PSD2024_Ricciato_v6_1.pdf and Belkin M., Hsu D., Ma S., & Mandal S. (2019), *Reconciling modern machine-learning practice and the classical bias–variance trade-off*. Proceedings of the National Academy of Sciences, 24 July 2019, 116(32) 15849-15854, available at: <https://www.pnas.org/doi/10.1073/pnas.1903070116>

³⁹ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024.

⁴⁰ Ibid, paragraph 1.

⁴¹ EDPB Report of the work undertaken by the ChatGPT Taskforce, adopted on 23 May 2024, paragraph 14.

*discriminatory, unexpected or misleading to the data subject*⁴². Considering the complexity of the technologies involved, information on the processing of personal data within AI models should therefore be provided in an accessible, understandable and user-friendly way⁴³. Transparency about the processing of personal data includes, in particular, compliance with the information obligations as set out in Articles 12 to 14 GDPR⁴⁴, which also require, in case of automated decision-making, including profiling, meaningful information about the logic involved, as well as the significance and the envisaged consequences of the processing for the data subject⁴⁵. Bearing in mind that the development phases of AI models may involve the collection of large amounts of data from publicly accessible sources (such as via web scraping techniques), reliance on the exception provided under Article 14(5)(b) GDPR is strictly limited to when the requirements of this provision are fully met⁴⁶.

64. **Purpose limitation and data minimisation principles** (Article 5(1)(b), (c) GDPR) - In accordance with the data minimisation principle, the development and deployment of AI models requires that personal data should be adequate, relevant and necessary in relation to the purpose. This can include the processing of personal data to avoid the risks of potential biases and errors when this is clearly and specifically identified within the purpose, and the personal data is necessary for that purpose (e.g. they cannot be effectively achieved by processing other data, including synthetic or anonymised data)⁴⁷. The WP29 already stressed that the *'purpose of the collection must be clearly and specifically identified [...]*⁴⁸. When assessing whether the purpose pursued is legitimate, specific and explicit, and whether the processing complies with the data minimisation principle, one should first identify the processing activity at stake. Notably, the different stages within the development or deployment phases may constitute the same or different processing activities, and may entail successive controllers or joint controllers. In some cases, it is possible to determine the purpose which will be pursued during the deployment of the AI model at an early development stage. Even where this is not the case, some context for that deployment should already be clear, and therefore, how this context informs the purpose of the development should be considered. When reviewing the purpose of the processing in a given stage of development, SAs should expect a certain degree of detail from the controller(s) and an explanation as to how these details inform the purpose of processing. This may include, for example, information on the type of AI model developed, its expected functionalities, and any other relevant context that is already known at that stage. The context of deployment could also include, for example, whether a model is being developed for internal deployment, whether the controller intends

⁴² EDPB Report of the work undertaken by the ChatGPT Taskforce, adopted 23 May 2024, paragraph 23; EDPB Guidelines 4/2019 on Article 25 Data Protection by Design and by Default, Version 2.0, adopted 20 October 2020, paragraph 69; Article 29 Working Party Guidelines on transparency under Regulation 2016/679, revised and adopted on 11 April 2018, endorsed by the EDPB on 25 May 2018, paragraph 2.

⁴³ Article 29 Working Party Guidelines on transparency under Regulation 2016/679, revised and adopted on 11 April 2018, endorsed by the EDPB on 25 May 2018, paragraph 5.

⁴⁴ See also Recital 39 GDPR, which states that it *'should be transparent to natural persons that personal data concerning them are collected, used, consulted or otherwise processed and to what extent the personal data are or will be processed [...]*.

⁴⁵ Article 13(2)(f) GDPR and Article 14(2)(g) GDPR.

⁴⁶ EDPB Report of the work undertaken by the ChatGPT Taskforce, adopted on 23 May 2024 paragraph 27.

⁴⁷ In addition, Article 10(5) AI Act provides for specific rules for the processing of special categories of personal data in relation to the high-risk AI systems for the purpose of ensuring bias detection and correction.

⁴⁸ Article 29 Working Party Opinion 03/2013 on purpose limitation (WP203), pp. 15-16.

to sell or distribute the model to third parties after its development, including whether the model is primarily intended to be deployed for research or commercial purposes.

65. **Data subject rights** (Chapter III GDPR) - Notwithstanding the need for SAs to ensure that all data subject rights are respected when AI models are developed and deployed by controllers, the EDPB recalls that whenever legitimate interest is relied upon as a legal basis by a controller, the right to object under Article 21 GDPR applies and should be ensured⁴⁹.

3.3.2 Considerations on the three steps of the legitimate interest assessment in the context of the development and deployment of AI models

66. In order to determine whether a given processing of personal data may be based on Article 6(1)(f) GDPR, SAs should verify that controllers have carefully assessed and documented whether the three following cumulative conditions are met: (i) the pursuit of a legitimate interest by the controller or by a third party; (ii) the processing is necessary to pursue the legitimate interest; and (iii) the legitimate interest is not overridden by the interests or fundamental rights and freedoms of the data subjects⁵⁰.

3.3.2.1 First step - Pursuit of a legitimate interest by the controller or by a third party

67. An interest is the broader stake or benefit that a controller or third party may have in engaging in a specific processing activity⁵¹. While the GDPR and the CJEU recognised several interests as being legitimate⁵², the assessment of the legitimacy of a given interest should be the result of a case-by-case analysis.
68. As recalled by the EDPB in its Guidelines on legitimate interest⁵³, an interest may be regarded as legitimate if the following three cumulative criteria are met:
- a. The interest is lawful⁵⁴;

⁴⁹ According to Article 21 GDPR, if a data subject objects, on grounds relating to their particular situation, to the processing of personal data concerning them, the controller shall no longer process the personal data, unless the controller demonstrates compelling legitimate grounds for the processing which override the interests, rights and freedoms of the data subject or for the establishment, exercise or defence of legal claims. Therefore, the two aspects to be taken into account by SAs is whether the controller is able to demonstrate such compelling overriding legitimate grounds and whether the right to object can be exercised.

⁵⁰ CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), para. 106; CJEU, judgment of 11 December 2019, Case C-708/18, *Asociația de Proprietari bloc M5A-ScaraA* (ECLI:EU:C:2019:1064), paragraph 40. See also EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0 adopted on 8 October 2024, paragraph 12 and ff. As reminded in these Guidelines, this 'assessment should be made at the outset of the processing, with the involvement of the Data Protection Officer (DPO) (if designated), and should be documented by the controller in line with the accountability principle set out in Article 5(2) GDPR'.

⁵¹ EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 14.

⁵² EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 16.

⁵³ EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 17.

⁵⁴ CJEU, judgment of 4 October 2024, Case C-621/22, *Koninklijke Nederlandse Lawn Tennisbond* (ECLI:EU:C:2024:857), paragraph 49, where the CJEU stressed that a legitimate interest cannot be contrary to the law. In this regard, the EDPB emphasises that, as the case may be, legislative frameworks should be taken into account when assessing the lawfulness of a given interest. See for example: Article 26(3) and Article 28 of Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market

- b. The interest is clearly and precisely articulated; and
 - c. The interest is real and present, not speculative.
69. Subject to the two other steps required by the legitimate interest assessment, the following examples may constitute a legitimate interest in the context of AI models: (i) developing the service of a conversational agent to assist users; (ii) developing an AI system to detect fraudulent content or behaviour; and (iii) improving threat detection in an information system.
- 3.3.2.2 Second step - Analysis of the necessity of the processing to pursue the legitimate interest*
70. The second step of the assessment consists in determining whether the processing of personal data is necessary for the purpose of the legitimate interest(s) pursued⁵⁵ (“necessity test”).
71. Recital 39 GDPR clarifies that ‘*Personal data should be processed only if the purpose of the processing could not reasonably be fulfilled by other means*’. According to the CJEU and previous EDPB guidance, the condition relating to the necessity of the processing should be examined in light of the fundamental rights and freedoms of the data subjects, and in conjunction with the data minimisation principle enshrined in Article 5(1)(c) GDPR⁵⁶.
72. The methodology referred to by the CJEU takes into account the context of the processing, as well as the effects on the controller and on the data subjects. The assessment of necessity therefore entails two elements: (i) whether the processing activity will allow the pursuit of the purpose⁵⁷; and (ii) whether there is no less intrusive way of pursuing this purpose⁵⁸.

For Digital Services and amending Directive 2000/31/EC (Digital Services Act) (‘DSA’) on forbidden targeted advertising to minors ; Article 5(1) and (2) of the AI Act on prohibited AI practices (manipulative practices and below the threshold of consciousness); processing in violation of intellectual property rights and the provisions in Directive (EU) 2019/790 on copyright and related rights in the Digital Single Market.

⁵⁵ EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraphs 28-30.

⁵⁶ CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), para. 108 and 109, also referring to CJEU, judgment of 11 December 2019, Case C-708/18, *Asociația de Proprietari bloc M5A-Scara A* (ECLI:EU:C:2019:1064), paragraph 48; CJEU, judgment of 9 November 2010, joined cases C-92/09 and C-93/09, *Volker und Markus Schecke* (ECLI:EU:C:2010:662), paragraphs 85 and 86; CJEU, judgment of 22 June 2021, Case C-439/19, *Latvijas Republikas Saeima* (ECLI:EU:C:2021:504), paragraphs 98, 109, 110, 113. See also for example: EDPB Guidelines 3/2019 on processing of personal data through video devices, version 2.0, adopted on 29 January 2020, paragraphs 24-26 and 73; EDPB Guidelines 2/2019 on the processing of personal data under Article 6(1)(b) GDPR in the context of the provision of online services to data subjects, version 2.0, adopted on 8 October 2019, paragraphs 23-25; EDPB Opinion 11/2024 on the use of facial recognition to streamline airport passengers’ flow, version 1.1, adopted on 23 May 2024, paragraph 27.

⁵⁷ See CJEU, judgment of 16 December 2008, Case C-524/06, *Heinz Huber v Bundesrepublik Deutschland* (ECLI:EU:C:2008:724), paragraph 66. Also in the same case, see the Opinion of Advocate General Poiares Maduro in Case C-524/06, *Heinz Huber v Bundesrepublik Deutschland* (ECLI:EU:C:2008:194), paragraph 16, stating: ‘*the proper test here is one of effectiveness, and it is for the national court to apply it. The question it must ask is whether there are other ways of data processing by which the immigration authorities could enforce the rules on residence status. If it answers that question in the affirmative, the centralised data storage and processing for Union citizens should be declared unlawful. It is not necessary for the alternative system to be the most effective or appropriate; it is enough for it to be able to perform adequately. Put differently, even if the central register is more effective or convenient or user-friendly than its alternatives (such as the decentralised, local registers), the latter are clearly to be preferred if they can be used to indicate the residence status of Union citizens*’.

⁵⁸ See CJEU, judgment of 27 September 2017, Case C-73/16, *Peter Puškár* (ECLI:EU:C:2017:725), paragraph 113: ‘*It is thus for the national court to ascertain whether the establishment of the contested list and the inclusion of the names of the data subjects in such a register are suitable for achieving the objectives pursued by them and*

73. For example, and as the case may be, the intended volume of personal data involved in the AI model needs to be assessed in light of less intrusive alternatives that may reasonably be available to achieve just as effectively the purpose of the legitimate interest pursued. If the pursuit of the purpose is also possible through an AI model that does not entail processing of personal data, then processing personal data should be considered as not necessary. This is particularly relevant for the development of AI models. When assessing whether the condition of necessity is met, SAs should pay particular attention to the amount of personal data processed and whether it is proportionate to pursue the legitimate interest at stake, also in light of the data minimisation principle.
74. The assessment of necessity should also take into account the broader context of the intended processing of personal data. The existence of means that are less intrusive to the fundamental rights and freedoms of the data subjects may vary depending on whether the controller has a direct relationship with the data subjects (first-party data) or not (third-party data). The CJEU provided some considerations to take into account when analysing the necessity of the processing of first-party data for the purpose of the legitimate interest(s) pursued (albeit in the context of disclosure of such data to third parties)⁵⁹.
75. Implementing technical safeguards to protect personal data may also contribute to meet the necessity test. This could include, for example, implementing measures such as those identified in Section 3.2.2 in such a way that anonymisation is not reached, but which still reduces the ease at which data subjects can be identified. The EDPB notes that some of these measures, when not required to comply with the GDPR, may constitute additional safeguards, as further analysed in the sub-section “mitigating measures”, within Section 3.3.2.3⁶⁰.

3.3.2.3 Third step - Balancing test

76. The third step of the legitimate interest assessment is the ‘**balancing exercise**’ (also referred to in this document as ‘**balancing test**’)⁶¹. This step consists in identifying and describing the different opposing rights and interests at stake⁶², i.e. on the one side the interests, fundamental rights and freedoms of the data subjects, and on the other side the interests of the controller or a third party. The specific circumstances of the case should then be considered to demonstrate that legitimate interest is an appropriate legal basis for the processing activities at stake⁶³.

whether there is no other less restrictive means in order to achieve those objectives’; See also for example Opinion of Advocate General Rantos in case C-252/21, Meta v. Bundeskartellamt, ECLI:EU:C:2022:704, paragraph 61, stating: ‘[...] It is necessary therefore for a close link to exist between the processing and the interest pursued, in the absence of alternatives that are more data-protection friendly, since it is not enough for the processing merely to be of use to the controller’.

⁵⁹ CJEU, judgment of 4 October 2024, Case C-621/22, *Koninklijke Nederlandse Lawn Tennisbond* (ECLI:EU:C:2024:857), paragraphs 51-53.

⁶⁰ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 57.

⁶¹ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraphs 31 to 60.

⁶² See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 32.

⁶³ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 32, also referring to CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 110.

Data subjects' interests, fundamental rights and freedoms

77. Article 6(1)(f) GDPR provides that, in assessing the different components in the context of the balancing test, the controller should take into account the interests, fundamental rights and freedoms of the data subjects. Data subjects' interests are those that may be affected by the processing at stake. In the context of the development phase of an AI model, these may include, but are not limited to, the interest in self-determination and retaining control over one's own personal data (e.g. the data gathered for the development of the model). In the context of the deployment of an AI model, the data subjects' interests may include, but are not limited to, interests in retaining control over one's own personal data (e.g. the data processed once the model is deployed), financial interests (e.g. where an AI model is used by the data subject to generate revenues, or is used by an individual in the context of their professional activity), personal benefits (e.g. where an AI model is used to improve accessibility to certain services), or socioeconomic interests (e.g. where an AI model enables access to better healthcare, or facilitates the exercise of a fundamental right such as access to education)⁶⁴.
78. The more precisely an interest is defined in light of the intended purpose of the processing, the better it will enable to clearly apprehend the reality of the benefits and risks to be taken into account in the balancing test.
79. In relation to the fundamental rights and freedoms of the data subjects, the development and deployment of AI models may raise serious risks to rights protected by the EU Charter of Fundamental Rights (the "EU Charter"), including, but not limited to, the right to private and family life (Article 7 EU Charter) and the right to protection of personal data (Article 8 EU Charter). These risks may occur during the development phase, for example where personal data is scraped against the data subjects' wishes or without their knowledge. These risks may also occur in the deployment phase, for example where personal data is processed by (or as part of) the model in a way which contravenes the data subjects' rights, or where it is possible to infer, accidentally or by attacks (e.g. membership inference, extraction, or model inversion), what personal data is contained in the learning database. Such situations present a risk for the privacy of data subjects whose data might appear in the deployment phase of the AI system (e.g. reputational risk, identity theft or fraud, security risk depending on the nature of the data).
80. Depending on the case at stake, there may also be risks to other fundamental rights. For example, large-scale and indiscriminate data collection by AI models in the development phase may create a sense of surveillance for data subjects, especially considering the difficulties to prevent public data from being scraped. This may lead individuals to self-censor, and present risks of undermining their freedom of expression (Article 11 EU Charter). In the deployment phase, risks for the freedom of expression are also present where AI models are used to block content publication from data subjects. In addition, an AI model recommending inappropriate content to vulnerable individuals may present risks for their mental health (Article 3(1) EU Charter). In other cases, the deployment of AI models may also lead to adverse consequences on the individual's right to engage into work (Article 15 EU Charter), for example when job applications are pre-selected using an AI model. In the same manner, an AI model could present risks for the right of non-discrimination (Article 21 EU Charter), if it discriminates individuals based on certain personal characteristics (such as nationality or gender). Furthermore, the

⁶⁴ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 38.

deployment of AI models may also present risks to the security and safety of the individual (e.g. where the AI model is used with malicious intent), as well as risks to their physical and mental integrity⁶⁵.

81. The deployment of AI models may also positively impact certain fundamental rights, e.g. the model may support the right to mental integrity of the person (Article 3 of the Charter), for instance when an AI model is used to identify harmful content online; or the model may facilitate the access to certain essential services or facilitate the exercise of fundamental rights, such as access to information (Article 11 EU Charter) or access to education (Article 14 EU Charter).

Impact of the processing on data subjects

82. The processing of personal data that takes place during the development and deployment of AI models may impact data subjects in different ways, which may be positive or negative⁶⁶. For example, if a processing activity entails benefits for the data subject, these may be taken into account in the balancing test. While the existence of such benefits may lead to the conclusion, by a SA, that the interests of the controller or a third party are not overridden by the interests, fundamental rights and freedoms of the data subjects, such conclusion may only be the result of a case-by-case analysis taking into consideration all appropriate factors.
83. The impact of the processing on the data subjects may be influenced by (i) the nature of the data processed by the models; (ii) the context of the processing; and (iii) the further consequences that the processing may have⁶⁷.
84. In relation to the **nature of the data processed**, it should be recalled that - apart from special categories of personal data and data relating to criminal convictions and offences that respectively enjoy additional protection under Articles 9 and 10 GDPR - the processing of some other categories of personal data may lead to significant consequences for data subjects. In this context, the processing of certain types of personal data revealing highly private information (e.g. financial data, or location data) for the development and deployment of an AI model should be considered as possibly having a serious impact on data subjects. In the deployment phase, consequences of such processing for data subjects may for example be economic (e.g. discrimination in the employment context) and/or reputational (e.g. defamation).
85. In relation to the **context of the processing**, it is first necessary to identify the elements that could create risks for the data subjects (e.g. the way in which the model was developed, the way in which the model may be deployed, and/or whether the security measures used to protect the personal data are appropriate). The nature of the model and the intended operational uses play a key role in identifying such potential causes.
86. It is also necessary to assess the severity of these risks for the data subjects. It may be considered, among other things, how the personal data is processed (e.g. if it is combined with other datasets), what the scale of the processing and the amount of personal data processed is⁶⁸ (e.g. the overall volume of data, the volume of data per data subject, the number of data subjects affected)⁶⁹, the

⁶⁵ Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 46.

⁶⁶ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 39.

⁶⁷ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 32.

⁶⁸ See EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 43.

⁶⁹ CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 116.

status of the data subject (e.g. children or other vulnerable data subjects) and their relationship with the controller (e.g. if the data subject is a customer). For example, the use of web scraping in the development phase may lead - in the absence of sufficient safeguards - to significant impacts on individuals, due to the large volume of data collected, the large number of data subjects, and the indiscriminate collection of personal data.

87. The **further consequences** that the processing may have should also be considered when assessing the impact of the processing on the data subjects. They should be assessed by SAs on a case-by-case basis, considering the specific facts at hand.
88. Such consequences may include (but are not limited to) risks of violation of the fundamental rights of the data subjects, as described in the previous sub-section⁷⁰. The risks may vary in likelihood and severity, and may result from personal data processing which could lead to physical, material or non-material damage, in particular where the processing may give rise to discrimination⁷¹.
89. Where the deployment of an AI model entails the processing of personal data of both (i) data subjects whose personal data is included in the dataset used in the development phase; and (ii) data subjects whose personal data is processed in the deployment phase, SAs should distinguish and consider the risks affecting the interests, rights and freedoms of each of these categories of data subjects when verifying the balancing test carried out by a controller.
90. **Lastly, the analysis of the possible further consequences of the processing should also consider the likelihood of these further consequences materialising.** The assessment of such likelihood should be made taking into consideration the technical and organisational measures in place and the specific circumstances of the case. For example, SAs may consider whether measures have been implemented to avoid a potential misuse of the AI model. For AI models which may be deployed for a variety of purposes, such as generative AI, this may include controls limiting as much as possible their use for harmful practices, for instance: the creation of deepfakes; chatbots that are used for disinformation, phishing and other types of fraud; and manipulative AI/Al agents (in particular where they are anthropomorphic or providing misleading information).

Reasonable expectations of the data subjects

91. Based on Recital 47 GDPR, *‘[a]t any rate the existence of a legitimate interest would need careful assessment including whether a data subject can reasonably expect at the time and in the context of the collection of the personal data that processing for that purpose may take place. The interests and fundamental rights of the data subject could in particular override the interest of the controller where personal data is processed in circumstances where data subjects do not reasonably expect further processing’*⁷².
92. Reasonable expectations play a key role in the balancing test, not least due to the complexity of the technology used in AI models and the fact that it may be difficult for data subjects to understand the

⁷⁰ See sub-section “Data subjects’ interests, fundamental rights and freedoms” above.

⁷¹ See Section 2.3 of the EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024. Also see recital 75 GDPR for further examples.

⁷² See also CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 112; CJEU, judgment of 11 December 2019, Case C-708/18, *Asociația de Proprietari bloc M5A-ScaraA* (ECLI:EU:C:2019:1064), paragraph 58; CJEU, judgment of 4 October 2024, Case C-621/22, *Koninklijke Nederlandse Lawn Tennisbond* (ECLI:EU:C:2024:857), paragraph 55.

variety of potential uses of an AI model and the data processing involved⁷³. To this end, the information provided to data subjects may be considered to assess whether data subjects can reasonably expect their personal data to be processed. However, while the omission of information can contribute to the data subjects not expecting a certain processing, the mere fulfilment of the transparency requirements set out in the GDPR is not sufficient in itself to consider that the data subjects can reasonably expect a certain processing⁷⁴. Further, simply because information relating to the development phase of an AI model is included in the controller's privacy policy, it does not necessarily mean that the data subjects can reasonably expect it to happen; rather, this should be analysed by SAs on the specific circumstances of the case and considering all of the relevant factors.

93. When assessing the reasonable expectations of data subjects in relation to processing that takes place in the development phase, it is important to refer to the elements mentioned in the EDPB Guidelines on legitimate interest⁷⁵. Further, within the subject-matter of this Opinion, it is important to consider the wider context of the processing. This may include, although is not limited to, whether or not the personal data was publicly available, the nature of the relationship between the data subject and the controller (and whether a link exists between the two), the nature of the service, the context in which the personal data was collected, the source from which the data was collected (e.g. the website or service where the personal data was collected and the privacy settings they offer), the potential further uses of the model, and whether data subjects are actually aware that their personal data is online at all.
94. In the development phase of the model, the data subjects' reasonable expectations may differ depending on whether the data processed to develop the model is made public by the data subjects or not. Further, the reasonable expectations may also differ depending on whether they directly provided the data to the controller (e.g. in the context of their use of the service), or if the controller obtained it from another source (e.g. via a third-party, or scraping). In both cases, the steps taken to inform the data subjects of the processing activities should be considered when assessing the reasonable expectations.
95. In the deployment phase of the AI model, it is equally important to consider the data subjects' reasonable expectations within the context of the model's specific capabilities. For example, for AI models which can adapt according to the inputs provided, it may be relevant to consider if the data subjects were aware that they had provided personal data so that the AI model could adjust its responses to their needs and so that they could obtain tailored services. Further, it may also be relevant to consider whether this processing activity would only impact the service provided to the data subjects (e.g. the personalisation of content for a specific user) or whether it would be used to modify the service provided to all customers (e.g. to improve the model in a general manner). As in the development stage, it may also be particularly relevant to consider whether there is a direct link between the data subjects and the controller. Such a direct link may, for example, allow the controller

⁷³ For example, in judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), paragraph 123, while the CJEU found that 'product improvement' cannot in principle be ruled out as a legitimate interest, it also found that it is *'doubtful whether [...] the 'product improvement' objective, given the scale of that processing and its significant impact on the user, as well as the fact that the user cannot reasonably expect those data to be processed [...] may override the interests and fundamental rights of such a user, particularly in the case where that user is a child'*.

⁷⁴ Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 53.

⁷⁵ Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraphs 50-54.

to easily provide information to the data subjects on the processing activity and the model, which could then influence those data subjects' reasonable expectations.

Mitigating measures

96. When the data subjects' interests, rights and freedoms seem to override the legitimate interest(s) being pursued by the controller or a third party, the controller may consider introducing mitigating measures to limit the impact of the processing on these data subjects. Mitigating measures are safeguards that should be tailored to the circumstances of the case and depend on different factors, including on the intended use of the AI model. These mitigating measures would aim to ensure that the interests of the controller or third party will not be overridden, so that the controller would be able to rely on this legal basis.
97. As recalled in the EDPB's Guidelines on legitimate interest, mitigating measures should not be confused with the measures that the controller is legally required to adopt anyway to ensure compliance with the GDPR, irrespective of whether the processing is based on Article 6(1)(f) GDPR⁷⁶. This is particularly important for measures that, for instance, require to comply with GDPR principles, such as the principle of data minimisation.
98. The list of measures provided below is non-exhaustive and non-prescriptive and the implementation of the measures should be considered on a case-by case-assessment. While, depending on the circumstances, some of the measures below may be required to comply with specific obligations of the GDPR, when this is not the case they may be taken into account as additional safeguards. In addition, some of the measures mentioned below relate to areas which are subject to rapid evolution and new developments, and should be taken into account by SAs when dealing with a specific case.
99. **In relation to the development phase of AI models**, several measures may be taken to mitigate risks posed by the processing of both first-party and third-party data (including to mitigate risks related to web scraping practices). On the basis of the above, the EDPB provides some examples of measures that may be implemented to mitigate the risks identified in the balancing test, and should be considered by SAs when assessing specific AI models on a case-by-case basis.
100. **Technical measures:**
 - a. Measures mentioned under Section 3.2.2 that are suitable to mitigate the risks at play, where those measures do not result in anonymisation of the model and are not required to comply with other GDPR obligations or under the necessity test (second step of the legitimate interest's assessment).
101. In addition to those, other relevant measures may include:
 - b. Pseudonymisation measures: this could, for example, include measures to prevent any combination of data based on individual identifiers. These measures may not be appropriate where the SA considers that the controller demonstrated the reasonable need to gather different data about a particular individual for the development of the AI system or model in question.
 - c. Measures to mask personal data or to substitute it with fake personal data in the training set (e.g. the replacement of names and email addresses with fake names and fake email

⁷⁶ Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 57.

addresses). This measure may be particularly appropriate when the actual substantive content of the data is not relevant to the overall processing (e.g. in LLM training).

102. **Measures that facilitate the exercise of individuals' rights:**

- a. Observing a reasonable period of time between the collection of a training dataset and its use. This additional safeguard may enable data subjects to exercise their rights during this period, with the reasonable period of time being assessed depending on the circumstances of each case.
- b. Proposing an unconditional 'opt-out' from the outset, for instance by providing a discretionary right to object to data subjects before the processing takes place, in order to strengthen the control of individuals over their data, which goes beyond the conditions of Article 21 GDPR⁷⁷.
- c. Allowing the data subjects to exercise their right to erasure even when the specific grounds listed in Article 17(1) GDPR do not apply⁷⁸.
- d. Allowing data subjects to submit claims of personal data regurgitation or memorisation and the circumstances and means by which the claim claims may be reproduced, allowing controllers to reproduce and assess relevant unlearning techniques to address the claims.

103. **Transparency measures:** in some cases, mitigating measures could include measures that provide for greater transparency with regard to the development of the AI model. Some measures, in addition to compliance with the GDPR obligations, may help overcoming the information asymmetry and allow data subjects to get a better understanding of the processing involved in the development phase:

- a. Release of public and easily accessible communications which go beyond the information required under Article 13 or 14 GDPR, for instance by providing additional details about the collection criteria and all datasets used, taking into account special protection for children and vulnerable persons.
- b. Alternative forms of informing data subjects, for instance: media campaigns with different media outlets to inform data subjects, information campaign by e-mail, use of graphic visualisation, frequently asked questions, transparency labels and model cards the systematisation of which could structure the presentation of information on AI models, and annual transparency reports on a voluntary basis.

104. **Specific mitigating measures in the context of web scraping:** Considering that, as mentioned above, web scraping raises specific risks⁷⁹, specific mitigating measures could be identified in this context. Where relevant, they may be considered by SAs, in addition to the mitigating measures mentioned above, when investigating controllers conducting web scraping.

105. Specific measures, when not necessary under the second step of the legitimate interest assessment, may prove useful to mitigate the risk in the context of web scraping. These may include **technical measures**, such as:

⁷⁷ Ibid.

⁷⁸ Ibid.

⁷⁹ These practices may also raise additional issues that are not covered under this Opinion, see for instance Pagallo U., Ciani Sciolla J., *Anatomy of web data scraping: ethics, standards, and the troubles of the law*. European Journal of Privacy Law & Technologies, (2023) 2 p. 1 - 19, available at: <https://doi.org/10.57230/EJPLT232PS>.

- a. Excluding data content from publications which might include personal data entailing risks for particular persons or groups of persons (e.g. individuals who might be subject to abuse, prejudice or even physical harm if the information were released publicly).
 - b. Ensuring that certain data categories are not collected or that certain sources are excluded from data collection; this could include, for instance, certain websites that are particularly intrusive due to the sensitivity of their subject matter.
 - c. Excluding collection from websites (or sections of websites) which clearly object to web scraping and the reuse of their content for the purpose of building AI training databases (for example, by respecting robots.txt or ai.txt files or any other recognised mechanism to express exclusion from automated crawling or scraping).
 - d. Imposing other relevant limits on collection, possibly including criteria based on time periods.
106. In the context of web scraping, examples of specific measures **facilitating the exercise of individuals' rights and transparency** may include: creating an opt-out list, managed by the controller and which allows data subjects to object to the collection of their data on certain websites or online platforms by providing information that identifies them on those websites, including before the data collection occurs⁸⁰.
107. **Specific considerations regarding mitigating measures in the deployment phase:** While some of the measures mentioned above may be also relevant for the deployment phase, depending on the circumstances, the EDPB provides below a non-exhaustive list of additional supporting measures that may be implemented and that should be assessed by SAs on a case-by-case basis.
- a. **Technical measures** may for instance be put in place to prevent the storage, regurgitation or generation of personal data, especially in the context of generative AI models (such as output filters), and/or to mitigate the risk of unlawful reuse by general purpose AI models (e.g. digital watermarking of AI-generated outputs).
 - b. **Measures that facilitate or accelerate the exercise of individuals' rights** in the deployment phase, beyond what is required by law, regarding in particular, and not limited to, the exercise of the right to erasure of personal data from model output data or deduplication, and post-training techniques that attempt to remove or suppress personal data.
108. When investigating the deployment of a specific AI model, SAs should consider whether the controller has published the balancing test it conducted, as it may increase transparency and fairness. As mentioned in the EDPB's Guidelines on legitimate interest, other measures may be considered to provide data subjects with information from the balancing test in advance of any collection of personal data⁸¹. The EDPB also reiterates⁸² that an element to be considered is whether the controller has involved the DPO, where applicable.

⁸⁰ Unless the controller demonstrates compelling legitimate grounds for the processing which override the interests, rights and freedoms of the data subject or for the establishment, exercise or defence of legal claims.

⁸¹ EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 68.

⁸² EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, Version 1.0, adopted on 8 October 2024, paragraph 12.

3.4 On the possible impact of an unlawful processing in the development of an AI model on the lawfulness of the subsequent processing or operation of the AI model

109. This section of the Opinion addresses Question 4 of the Request. This Question seeks clarification on the possible impact of an unlawful processing in the development phase on the subsequent processing (for instance in the deployment phase of the AI model) or on the operation of the model. The question seeks to address both the situation where such an AI model processes personal data which is retained in the model (Question 4(i) of the Request), as well as the situation where no personal data processing is involved anymore in the deployment of the AI model (i.e. the model is anonymous) (Question 4(ii) of the Request).
110. Before addressing certain specific scenarios, the EDPB provides the following general considerations.
111. First, the clarifications provided in this section will focus on the processing of personal data in the development phase conducted while non-complying with the principle of lawfulness as set out in Article 5(1)(a) GDPR and Article 6 GDPR specifically (hereafter “**unlawfulness**”)⁸³. In the same vein, the considerations of the EDPB will focus on the impact of the unlawfulness of the processing in the development phase on the lawfulness (i.e. compliance with Article 5(1)(a) GDPR and Article 6 GDPR) of the subsequent processing or operation of the model. However, the EDPB points out that the processing conducted in the development phase may also lead to infringements of other GDPR provisions, such as the lack of transparency towards data subjects, or data protection by design and/or default, which are not analysed in this Opinion.
112. Second, when addressing this question, the accountability principle, which requires controllers to be responsible for, and demonstrate compliance with, *inter alia*, Article 5(1) GDPR and Article 6 GDPR⁸⁴, plays a key role. This is also true for the need to assess which organisation is a controller for the processing activity at stake, and whether situations of joint-controllership arise (as they may be inextricably linked)⁸⁵. Considering the significance of the factual circumstances of each case, including with regard to the role played by each party involved in the processing, the considerations of the EDPB should be understood as general observations which should be assessed on a case-by-case basis by SAs.
113. Third, the EDPB highlights that, in accordance with Article 51(1) GDPR, SAs are ‘*responsible for monitoring the application of [the GDPR], in order to protect the fundamental rights and freedoms of natural persons in relation to processing and to facilitate the free flow of personal data within the Union*’. It is therefore within the SAs’ competence to assess the lawfulness of the processing and to exercise their powers granted by the GDPR in line with their national framework⁸⁶. In such cases, SAs enjoy discretionary powers to assess the possible infringement(s) and choose appropriate, necessary

⁸³ CJEU, judgment of 4 May 2023, Case C-60/22, *Bundesrepublik Deutschland* (ECLI:EU:C:2023:373), paragraphs 55-57.

⁸⁴ CJEU, judgment of 4 May 2023, Case C-60/22, *Bundesrepublik Deutschland* (ECLI:EU:C:2023:373), paragraph 53.

⁸⁵ EDPB Guidelines 07/2020 on the concepts of controller and processor in the GDPR, Version 2.1, adopted on 7 July 2021, paragraph 55.

⁸⁶ Specific national rules may have to be taken into account. See for instance Article 2-decies of the Italian Data Protection Code (legislative decree 196/2003) which establishes that the data processed in violation of the data protection rules cannot be used. This is without prejudice to other national legal frameworks, such as criminal laws.

and proportionate measures, among those mentioned under Article 58 GDPR, taking into account the circumstances of each individual case⁸⁷.

114. **When there is the finding of an infringement, SAs may impose corrective measures, such as ordering controllers, taking into account the circumstances of each case, to take actions in order to remediate the unlawfulness of the initial processing.** These may include, for instance, issuing a fine, imposing a temporary limitation on the processing, erasing part of the dataset that was processed unlawfully or, where this is not possible, depending on the facts at hand, having regard to the proportionality of the measure, ordering the erasure of the whole dataset used to develop the AI model and/or the AI model itself. When assessing the proportionality of the envisaged measure, SAs may take into account measures that can be applied by the controller to remediate the unlawfulness of the initial processing (e.g. retraining).
115. The EDPB also highlights that, when personal data is processed unlawfully, data subjects can request deletion of their personal data, subject to the conditions set forth under Article 17 GDPR, and that SAs may order the erasure of the personal data *ex officio*⁸⁸.
116. When assessing whether a measure is appropriate, necessary, and proportionate, SAs may consider, among other elements, the risks raised for the data subjects, the gravity of the infringement, the technical and financial feasibility of the measure, as well as the volume of personal data involved.
117. Finally, the EDPB recalls that the measures taken by SAs under the GDPR are without prejudice to those taken by competent authorities under the AI Act and/or under other applicable legal frameworks (e.g. legislation on civil liability).
118. In the subsequent sections, the EDPB will address three scenarios covered by Question 4 of the Request, where the differences lie on whether the personal data processed to develop the model is retained in the model, and/or whether the subsequent processing is performed by the same or another controller.

3.4.1 Scenario 1. A controller unlawfully processes personal data to develop the model, the personal data is retained in the model and is subsequently processed by the same controller (for instance in the context of the deployment of the model)

119. This scenario relates to Question 4(i) of the Request, in the situation where a controller unlawfully processes personal data (i.e. by non-complying with Article 5(1)(a) GDPR and Article 6 GDPR) to develop an AI model, the AI model retains information relating to an identified or identifiable natural person and thus is not anonymous. Personal data is then subsequently processed by the same controller (for instance in the context of the deployment of the model). With regard to this scenario, the EDPB provides the following considerations.

⁸⁷ See in this regard, Recital 129 GDPR, as well as CJEU, judgment of 26 September 2024, Case C-768-21, *TR v Land Hessen* (ECLI:EU:C:2024:785), paragraph 37; CJEU judgment of 7 December 2023, in Joined Cases C-26/22 and C-64/22, *SCHUFA Holding (Libération de reliquat de dette)* (ECLI:EU:C:2023:958), paragraph 57; and CJEU, judgment of 14 March 2024, Case C-46/23, *Újpesti Polgármesteri Hivatal* (ECLI:EU:C:2024:239), paragraph 34.

⁸⁸ In this regard, EDPB Opinion 39/2021 on whether Article 58(2)(g) GDPR could serve as a legal basis for a supervisory authority to order *ex officio* the erasure of personal data in a situation where such request was not submitted by the data subject, paragraph 28. See also, in this regard, CJEU, judgment of 14 March 2024, Case C-46/23, *Újpesti Polgármesteri Hivatal* (ECLI:EU:C:2024:239), paragraph 42.

120. The power of the SA to impose corrective measures on the initial processing (as explained under paragraphs 113, 114, 115 above), would in principle have an impact on the subsequent processing (e.g. if the SA orders the controller to delete the personal data that was processed unlawfully, such corrective measures would not allow the latter to subsequently process the personal data that was subject to the measures).
121. With specific regard to the impact of the unlawful processing in the development phase on the subsequent processing (for instance in the deployment phase), the EDPB recalls that it is for SAs to conduct a case-by-case analysis that takes into account the specific circumstances of each case.
122. **Whether the development and deployment phases involve separate purposes (thus constituting separate processing activities) and the extent to which the lack of legal basis for the initial processing activity impacts the lawfulness of the subsequent processing, should be assessed on a case-by-case basis, depending on the context of the case.**
123. For instance, with specific regard to the legal basis of Article 6(1)(f) GDPR, when the subsequent processing is based on legitimate interest, the fact that the initial processing was unlawful should be taken into account in the legitimate interest assessment (e.g. with regard to the risks for data subjects or the fact that data subjects may not expect such subsequent processing). In these cases, the unlawfulness of the processing in the development phase may impact the lawfulness of the subsequent processing.

3.4.2 Scenario 2. A controller unlawfully processes personal data to develop the model, the personal data is retained in the model and is processed by another controller in the context of the deployment of the model

124. This scenario relates to Question 4(i) of the Request. It differs from scenario 1 (in Section 3.4.1 of this Opinion) as personal data is subsequently processed by another controller in the context of the deployment of the AI model.
125. The EDPB recalls that ascertaining the roles assigned to these different actors under the data protection framework is an essential step in order to identify which obligations under the GDPR apply and who is responsible for those obligations, and that joint controllership situations should also be considered when assessing each parties' responsibilities under the GDPR. Therefore, the observations below should be considered as general elements that should be taken into account by SAs where relevant. With regard to this scenario 2, the EDPB provides the following considerations.
126. First, it should be recalled that, according to Article 5(1)(a) GDPR, read in light of Article 5(2) GDPR, each controller should ensure the lawfulness of the processing it conducts and be able to demonstrate it. Therefore, SAs should assess the lawfulness of the processing carried out by (i) the controller that originally developed the AI model; and (ii) the controller that acquired the AI model and processes the personal data by itself.
127. Second, the consideration made under paragraphs 113, 114, 115 above are relevant in this case, with regard to the power of SAs to intervene in relation to the initial processing. Article 17(1)(d) GDPR (erasure of unlawfully processed data) and Article 19 GDPR (notification obligation regarding rectification or erasure of personal data or restriction of processing) may, depending on the circumstances of the case, also be relevant in this context, for instance in relation to the notification that the controller developing the model should conduct towards the controller deploying the model.

128. Third, in relation to the possible impact of the unlawfulness of the initial processing on the subsequent one conducted by another controller, such assessment should be conducted by SAs on a case-by-case basis.
129. **SAs should take into account whether the controller deploying the model conducted an appropriate assessment, as part of its accountability obligations⁸⁹ to demonstrate compliance with Article 5(1)(a) and Article 6 GDPR, to ascertain that the AI model was not developed by unlawfully processing personal data.** Such evaluation by SAs should take into account whether the controller has assessed some non-exhaustive criteria, such as the source of the data and whether the AI model is the result of an infringement of the GDPR, particularly if it was determined by a SA or a court, so that the controller deploying the model could not ignore that the initial processing was unlawful.
130. The controller should consider, for instance, if the data originates from a personal data breach or if the processing was subject to the finding of an infringement from a SA or a court. **The degree of the assessment of the controller and the level of detail expected by SAs may vary depending on diverse factors, including the type and degree of risks raised by the processing in the AI model during its deployment in relation to the data subjects whose data was used to develop the model.**
131. The EDPB notes that the AI Act requires providers of high-risk AI systems to draw up an EU declaration of conformity⁹⁰, and that such declaration contains a statement that the relevant AI system complies with EU data protection laws⁹¹. The EDPB notes that such a self-declaration may not constitute a conclusive finding of compliance under the GDPR. It may nonetheless be taken into account by the SAs when investigating a specific AI model.
132. The same considerations made under paragraph 123 above are also relevant in this case. When SAs verify if and how the controller assessed the appropriateness of legitimate interest as a legal basis for the processing it conducts, the unlawfulness of the initial processing should be taken into account as part of the legitimate interest assessment, for instance by assessing the potential risks that may arise for the data subjects whose personal were unlawfully processed to develop the model. Different aspects, either of a technical nature (e.g. the existence of filters or access limitations placed during the development of the model, which the subsequent controller cannot circumvent or influence, and which might prevent access to or disclosure of personal data) or of a legal nature (e.g. nature and severity of the unlawfulness of the initial processing) have to be given due consideration within the balancing test.

3.4.3 Scenario 3. A controller unlawfully processes personal data to develop the model, then ensures that the model is anonymised, before the same or another controller initiates another processing of personal data in the context of the deployment

133. This scenario relates to Question 4(ii) of the Request and refers to a case where a controller unlawfully processes personal data to develop the AI model, but does so in a way that ensures that personal data is anonymised, before the same or another controller initiates another processing of personal data in the context of the deployment. First, the EDPB recalls that SAs are competent and have the power to intervene with regard to the processing related to the anonymisation of the model, as well as to the processing conducted during the development phase. Thus, SAs may, depending on the specific

⁸⁹ Article 5(2) GDPR and Article 24 GDPR.

⁹⁰ Article 16(g) and Article 47 AI Act.

⁹¹ Annex V, point 5 AI Act.

circumstances of the case, impose corrective measures on this initial processing (as explained under paragraphs 113, 114, 115 above)

134. If it can be demonstrated that the subsequent operation of the AI model does not entail the processing of personal data, the EDPB considers that the GDPR would not apply⁹². Hence, the unlawfulness of the initial processing should not, therefore, impact the subsequent operation of the model. However, the EDPB emphasises that a mere assertion of anonymity of the model is not enough to exempt it from the application of the GDPR, and notes that SAs should assess it taking into account, on a case-by-case basis, the considerations provided by the EDPB to address Question 1 of the Request.
135. **When the controllers subsequently process personal data collected during the deployment phase, after the model has been anonymised, the GDPR would apply in relation to these processing activities. In these cases, as regards the GDPR, the lawfulness of the processing carried out in the deployment phase should not be impacted by the unlawfulness of the initial processing.**

4 Final remarks

136. This Opinion is addressed to all SAs and will be made public pursuant to Article 64(5)(b) GDPR.

For the European Data Protection Board

The Chair

Anu Talus

⁹² Recital 26 GDPR.