LAWGENT

The (ab)use of personal and supposedly anonymised data to train AI models including General-Purpose AI (GPAI) models:
A warning

Douwe Korff
Emeritus Professor of International Law, London Metropolitan University
Associate, Oxford Martin School, University of Oxford

Prepared for Lunch and Learn with Douwe Korff, 23 October 2026
Lawgent AB · Building the Future of Legal Advisory in Europe
October 2026

1. INTRODUCTION

1.1 Artificial Intelligence (AI)

The Stanford Institute for Human-Centered AI (HAI) describes Artificial Intelligence (AI) as follows:[1]

> **Artificial Intelligence (AI)** is a term coined in 1955 by John McCarthy, Stanford's first faculty member in AI, who described it as "the science and engineering of making intelligent machines." Today it is a broad term for **computer systems that can perform tasks with human-like intelligence, such as understanding language, recognizing images, learning from data, reasoning, and making decisions. Modern AI often works by finding patterns in large amounts of data and using those patterns to generate predictions or responses.** It can be narrow (good at a specific task) or more general-purpose, like today's large language models that can handle many tasks.

Although there are a range of different AI models,[2] in this paper I will focus on AI models that are "trained" on very large amounts of personal data (often scraped from Internet websites such as Spotify or records on hundreds of millions of airline passengers) to find patterns (or correlations) and that then suggest (or autonomously decide on) actions in relation to singled-out individuals. Typically, this is done by the model creating a profile or score from a cluster of attributes that it finds correlate to something (like creditworthiness, productivity or a likelihood of [re-]offending, or "possibly being a terrorist") and then singling out individuals to which the profile or score applies. Increasingly, people are granted or denied credit, or a job, or early release from prison, or placed on a "no fly list", etc. (or even targeted for assassination) on the basis of such AI-generated suggestions – either because a person or persons (bank official, admissions officer, parole board, soldier) relies on the score/suggestion, or because the model is allowed to take the relevant decision itself.

Many believe AI is heralding a "new industrial revolution" that will transform the world as we know it[3] – but perhaps not for the better: recently, there have been warnings that AI poses an "existential threat" to humanity[4] and "could cause a billion deaths" if it fell into the wrong hands.[5]

1.2 The EU "release the potential of data" Digital Strategy

In spite of the apocalyptic warnings, the EU Commission is pushing ahead with an EU "Data Union Strategy" that is aimed at "Unlocking Data for AI".[6] More specifically, on 19 November 2025, the European Commission released its so-called "Digital Omnibus Package".[7] At its core

is the Digital Omnibus, a proposed EU regulation,[8] accompanied by a 205-page Staff Working Document.[9] The Digital Omnibus seeks to:[10]

> Unlock[ ] opportunities in the use of data, as a fundamental resource in the EU economy, not least in view of supporting the development and use of trustworthy artificial intelligence solutions in the EU market.

In the Digital Omnibus, the Commission proposes what it calls "*targeted amendments to the [EU] data protection and privacy rules*" (and, as noted later, the EU has already adopted amendments to the EU AI Act that only recently started getting into application) to "*support this objective and provide immediate simplification measures for businesses and individuals, strengthening their ability to exercise their rights.*"[11]

In the next section, I briefly discuss (with references in endnotes to earlier papers that expand on the issues): first, the main EU data protection and rule-of-law guarantees relevant to the training and use of AI. Then, second, the relevant rules in the EU AI Act (and important, problematic exclusions of certain AI models and -systems[12] from that Act) and the interactions between the rules in the GDPR and the AI Act. These show the current EU "guardrails" for AI (and the gaps in them).

In section 3, I then discuss the proposed changes (insofar as these are becoming clear) and a related consultation on guidelines on anonymisation, and the possible implications.

In section 4, I conclude that if the changes and guidelines are adopted as proposed they would drive a horse and cart through the existing rules with a clear risk of untrammeled use of personal and not-really-anonymised data to train increasingly autonomous and dangerous AI systems and -models including especially General-Purpose AI models (GPAI).

2. THE CURRENT EU GUARDRAILS FOR AI

2.1 The EU data protection guarantees and rules

Article 8 of the EU Charter of Fundamental Rights stipulates the following:

> *Article 8*
>
> **Protection of personal data**
>
> 1. Everyone has the right to the protection of personal data concerning him or her.
> 2. Such data must be processed fairly for specified purposes and on the basis of the consent of the person concerned or some other legitimate basis laid down by law. Everyone has the right

of access to data which has been collected concerning him or her, and the right to have it rectified.

3. Compliance with these rules shall be subject to control by an independent authority

This right is further elaborated on in the EU General Data Protection Regulation (GDPR),[13] the somewhat orphaned e-Privacy Directive,[14] and the GDPR's sister instruments, the Law Enforcement Directive[15] and the regulation regulating the use of personal data by the EU institutions themselves, Regulation 2018/1725.[16] I will limit myself here to the GDPR.

Seven GDPR (and CFR-related) issues are crucial to the use of personal data to train and use AI, in particular in profiling and automated decision-making:

- the definitions and scope of the concepts of "personal data", "pseudonymisation" and "anonymisation": (subject to an important *caveat* to which I will come at *v*, below, and 3.2, below) if certain data fall within the concept of "personal data", they are subject to the constitutional protection accorded by Article 8 CFR and to the specific rules in the GDPR; if not, they are not protected by the CFR or the GDPR at all;
- the stipulations in both Article 8 CFR and Article 5 GDPR that personal data must be processed "fairly" and transparently;
- the further stipulation in these same articles that personal data may only be processed for "specified purposes" and "compatible" ("not incompatible") purposes; and the special rules on presumably compatible processing for "research" purposes (provided certain important conditions are met and certain safeguards are put in place);
- the requirement in Article 8 CFR that processing of personal data must have a "legitimate basis" and the further details on what constitute such bases in Article 6 GDPR, and on the additional conditions in Article 9 that must be met for processing of "sensitive data" (data on race, religion, health, sexual orientation, etc.);
- the in-principle prohibition in Article 22 GDPR of "profiling" and the taking of fully-automated decisions on individuals (subject to limited, conditional exceptions) and the associated rights of data- and decision subjects to be provided with meaningful information on the "logic" used in such decisions and to challenge such decisions including the right to obtain (meaningful) "human intervention";

- the fact that the very act of "anonymising" personal data constitutes "processing" of those data and must therefore comply with all the requirements relating to processing of personal data in the GDPR; and
- the obligation to carry out a Data Protection Impact Assessment (DPIA) in certain cases including in relation to automated decision-making.

i. "Personal data"[17]

"Personal data" are defined in Article 3(1) GDPR as follows:

> '**personal data**' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.

Recital 26 adds the following:

> The principles of data protection should apply to any information concerning an identified or identifiable natural person. Personal data which have undergone pseudonymisation, which could be attributed to a natural person by the use of additional information should be considered to be information on an identifiable natural person. To determine whether a natural person is identifiable, account should be taken of all the means reasonably likely to be used, such as **singling out**, either by the controller or by another person to identify the natural person directly or indirectly. To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments. The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable. This Regulation does not therefore concern the processing of such anonymous information, including for statistical or research purposes.
>
> (emphasis added)

The question of "singling out" a person from others in a wider context or a dataset is crucial: if a person can be "singled out" from a dataset, even without being named, and different data ("records") relating to that singled-out person can be linked, so that conclusions can be inferred from the links, that person is identifiable within the meaning of the GDPR, and those data are therefore personal data.[18]

Recital 26 says that "pseudonymous data" must always be regarded as personal data, but that, conversely, "*the principles of data protection [do not] apply [at all] to anonymous data*". But

recent case-law and guidelines refine this. Specifically, in its *SRB* judgment, the Court of Justice of the EU held that: [19]

> [P]seudonymised data must not be regarded as constituting, in all cases and for every person, personal data for the purposes of the application of [EU data protection law], in so far as pseudonymisation may, depending on the circumstances of the case, effectively prevent persons other than the controller from identifying the data subject in such a way that, for them, the data subject is not or is no longer identifiable.

In other words, if an entity is given (or otherwise obtains, perhaps even unlawfully) access to pseudonymous data subject to measures that really, actually and effectively prevent that entity from identifying the data subjects (by name or by otherwise singling out specific persons), then those data do not constitute personal data for that entity.[20] Conversely, if, in a specific context, the measures taken do not really, actually and effectively prevent re-identification (in either sense), then the relevant pseudonymous data must be regarded as still constituting personal data for any entity that can overcome those measures.

Indeed, as the Court held in its *Gesamtverband Autoteilehandel* judgment, data which are in themselves impersonal may become "personal" in nature where the controller puts them at the disposal of other persons who have means reasonably likely to enable the data subjects to be identified.[21]

Crucially, in the special circumstances of the *SRB* case, the recipient of the data was actually, really prevented from identifying the data subjects. In most cases of disclosure of pseudonymised data to third parties that is not the case. On the contrary: usually, data are provided to third parties in pseudonymised format precisely in order to allow the recipient to link distinct records relating to a specific data subject to draw conclusions from those links. For instance, a hospital may make patient records available to researchers in pseudonymised format so that those researchers can identify if certain complications arise in certain patients in certain situations, e.g., when certain medicines are used. The discovering of such correlations is of course also precisely what artificial intelligence models are trained to do.

However, as I will note below, at 3, the European Commission tried – in my opinion, wrongly and indeed disingenuously – to rely on the *SRB* judgment to open up wider, less constrained use of pseudonymous data, in particular also in the context of AI; and the guidelines on anonymisation, recently issued for consultation by the European Data Protection Board, in my opinion also leave too much wriggle-room in that respect.

ii. "Fair and transparent" processing

There are two elements to the Charter requirement that personal data should be processed "fairly" that are of special importance to the use of personal data in artificial intelligence.

First of all, as Article 5(1)(a) GDPR stipulates, personal data must be processed "fairly ***and in a transparent manner in relation to the data subject***": transparency *viz-à-viz* data subjects is an essential aspect of fairness. If you do not know what personal data are being collected and used in relation to you, for what purposes, and with whom the data are shared, you are denied agency ("informational self-determination") in relation to the processing. If you do not know these matters, you become a helpless object, incapable of exercising your rights – including your right under Article 47 of the EU Charter of Fundamental Rights to an "effective remedy" against matters that affect your fundamental rights and freedoms.

Articles 12 – 14 GDPR consequently provide for considerable detail on how the required transparency must be provided for. Crucially, for the present context, this includes the stipulation that (both when data are collected directly from the data subject, and when they are collected otherwise), the data subject must be provided with:

> [information about] the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, **meaningful information about the logic involved**, as well as the **significance** and the envisaged **consequences** of such processing for the data subject.
>
> (Article 13(2)(f) and 14(2)(g), emphasis added)

I have discussed the important requirements that flow from this obligation in some detail in a previous paper.[22] Here, it must suffice to reiterate the main points and conclusions.

Specifically, I noted that in its *PNR* judgment, the Court effectively prohibited the use of certain self- or machine-learning (i.e., self-modifying) AI systems ("ML/AI systems") – what an UK Information Commissioner's Office/Alan Turing Institute paper referenced in that previous paper calls "**black box**" systems[23] – in the taking of decisions under that directive, in terms that have wider implications:[24]

> [T]he use of artificial intelligence technology in self-learning systems ('machine learning') [that are] capable of modifying without human intervention or review the assessment process and, in particular, the assessment criteria on which the result of the application of that process is based … would be liable to render redundant the individual review of [the relevant decisions]. … Given the opacity which characterises the way in which artificial intelligence technology works, it might be impossible to understand the reason why a given program arrived at a positive match. In those circumstances, use of such technology may deprive the data subjects also of their right to an effective judicial remedy enshrined in Article 47 of the

> Charter, for which the PNR Directive, according to recital 28 thereof, seeks to ensure a high level of protection, in particular in order to challenge the non-discriminatory nature of the results obtained.

As I concluded in another opinion:[25]

> **It is clear from the PNR judgment that the use of opaque, unreviewable and unchallengeable "black box" ML/AI systems in the taking of any "significant" decisions on individuals is inherently incompatible with the EU Charter of Fundamental Rights. In particular, the use of such systems violates the very essence both of the right to respect for one's personality and of the right to an effective remedy.**

Secondly, any processing of personal data that results in **discrimination** against protected groups is ***unfair*** and therefore in violation of the Charter and the GDPR. That is the case irrespective of whether the discriminatory effects were intended or not. In another earlier paper, I noted that while Article 21 of the Charter of Fundamental Rights (and a host of global and European non-discrimination treaties and instruments)[26] and the GDPR require that all such discrimination must be "eliminated" (and prevented), the AI Act and the Commission Digital Omnibus proposals suggest that it suffices if biases in the outputs of AI models and -systems are "minimised".[27]

Here, it must suffice to reproduce the summary and conclusions at the end of that paper:

> The title of this paper put a (not so simple) question: **must all algorithmic discrimination be "prevented" and "eliminated" – or need it only be "minimised"?** My analysis shows that the GDPR supports the strong international human rights prohibition of discrimination: recital 71 stresses that all discriminatory outcomes of automated decision-making and profiling must be "prevented".
> In its *SCHUFA* judgment, the EU Court of Justice has further clarified that, in relation to any decisions that are based "decisively" on scores created by AI, the creation of the score already constitutes an "automated decision" within the meaning of Article 22(1) GDPR – and it follows from recital 71 that therefore any discriminatory effects of the score must be "prevented". Merely trying to "minimise" such effects is not good enough.
> However, the complete avoidance of bias in the datasets used in AI systems to generate outputs such as scores ("profiles" within the meaning of the GDPR) is in many cases effectively impossible – which is why the more recent EU rules on digital matters merely require that such biases in outputs be "minimised".
> The [UK Information Commissioner's Office] extrapolates from this to suggest that the effects of the outputs – in particular, any discriminatory outcomes of the use of the outputs – also only need to be "mitigated" (rather than altogether "prevented"). It wrongly asserts that:[28]
> It is unrealistic to adopt a 'zero tolerance' approach to risks to rights and freedoms [of individuals], and indeed the law does not require you to do so.

> This is likely to also be a view taken by "Big Tech" companies developing AI systems. And indeed, perhaps, also by the European Commission.

In fact, the law – more specifically, a massive corpus of international anti-discrimination law[29] – does precisely that: like the GDPR, it requires a "zero tolerance" approach to discrimination.

> The European Commission suggests that "keeping humans in the loop" somehow can remedy biased outputs from AI systems. However, given that such scores are typically created in "black box" AI systems, it will be well-nigh impossible – and for third parties relying on the score without insight into the algorithms used, effectively completely impossible – to "mitigate" the discriminatory effects (outcomes).
> There is a serious dilemma here: should we follow the suggestions of the UK ICO and the Alan Turing Institute and accept a certain level of discriminatory outcomes from the use of (in particular, "black box") AI systems? Or should we maintain the old concensus that all discrimination on the basis of race, religion, beliefs, health, etc., should be eliminated? Even if this means foregoing the use of AI systems that could bring significant societal benefits (e.g., in "identifying" possible/probable fraudsters)?
> ***As a human rights lawyer, I do not want to see a society in which opaque systems wrongly and disproprotionally "identify" more black people, or immigrants, or people with handicaps, as "probable" criminals. We should uphold the most basic principles of human rights law – not sell them out by rejecting them as "unrealistic".***
> (original emphases)

I will return to this at *v.*, below, in relation to automated decision-making.

iii. Purpose-specification and -limitation

As noted above, Article 8(2) CFR stipulates that "*[personal data] must be processed … for **specified purposes***". Article 5(1)(b) somewhat qualifies this by stipulating that:

> [Personal data shall be] collected for **specified, explicit and legitimate purposes** and *not further processed in a manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, in accordance with Article 89(1), not be considered to be incompatible with the initial purposes* ('purpose limitation').

I will come to the issue of "compatible purposes" later. Here, I should first of all point out that, as the Article 29 Working Party put it, "*purpose specification lies at the core of the legal framework established for the protection of personal data*".[30] In particular, the associated data protection requirements, that personal data must be "*adequate, relevant and limited to what is necessary*" (Art. 5(1)(c) GDPR) can only be met if the purpose(s) of the processing is (are) known (and "specified"): the adequacy, relevance and necessity of any personal data that are collected and processed can only be determined by reference to the "specified purpose[s]" for which they are collected. And this in turn means that the purpose(s) in question must be specified in sufficient detail:[31]

> **The purpose of the [data] collection must be clearly and specifically identified**: it must be detailed enough to determine what kind of processing is and is not included within the specified purpose, and to allow that compliance with the law can be assessed and data protection safeguards applied.
>
> For these reasons, **a purpose that is vague or general**, such as for instance 'improving users' experience', 'marketing purposes', 'IT-security purposes' or 'future research' will - without more detail - **usually not meet the criteria of being 'specific'**. That said, the degree of detail in which a purpose should be specified depends on the particular context in which the data are collected and the personal data involved. In some clear cases, simple language will be sufficient to provide appropriate specification, while in other cases more detail may be required.

If data are to be used for a range of purposes, "*sufficient granularity will be needed to make sure that all the different purposes are sufficiently clear [i.e., sufficiently clearly specified] for the [data subjects] [to understand]*".[32]

The issue is important here because **it has been suggested that "training" of an AI model (in the first phase of a two-phase process) can be seen as a separate purpose from the purpose for which the model is to be used (the second phase of the process)** and that the processing in each phase can be based on different legal bases.

I have criticised this in two blogs.[33] There, I explain in more detail than is possible here that, **in my opinion, in relation to purpose-specific, "narrow" AI models, this splitting up of purposes makes no sense.** If an AI system is being developed for, say, the purpose of direct marketing, or to improve medical treatment of patients with a particular illness (or to improve the diagnosis of a particular illness), then surely the training of the system also serves that purpose: the training of the system cannot be separated from the ultimate purpose of the system; training such an AI system is not a purpose in itself. This has a major bearing on the issue of compatibility, as noted below and (as already indicated) on the question of whether the "training" of the models can be based on a different legal basis than the processing for which the model is used (on which, see sub-section *iv*, below).

Even more profound issues arise in relation to so-called "General-Purpose AI" (GPAI) such as GPT-4, Claude and Gemini. This is defined in the AI Act as: [34]

> "an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications" (Article 3(63)).

The core point to note here is that GPAI[35] can serve "a wide range of tasks", i.e., tasks performed (or to be performed) for a wide range of purposes that are not specified (defined or delineated) from the start. In my opinion, this raises questions of whether such GPAI models are fundamentally compatible with the purpose-specification and -limitation principle.

In my opinion, General Purpose AI is *by its very nature* incompatible with the purpose-specification principle.

More specifically, it is difficult to see how it can be ensured that the "large amounts of data" on which they are "trained" – and which are often "scraped" from ubiquitous sources such as the Internet or massive data collections ("Big Data") – are limited to what is "necessary", "relevant", etc., to the *unspecified* (indeed, unlimited) range of purposes that the model may serve. Unless, of course, again, one were to regard the "training" – indeed, the "self-training" – of the model as a purpose in its own right.

But to define the purpose of GPAI model training as "developing a self-learning, self-amending algorithm that can serve any purpose" rather robs the concept of purpose-specification of any meaning. More specifically, if (as the WP29 and the EDPB confirm) the whole point of the principle is to allow for a test of what data are necessary and relevant in relation to the "*specified* purpose", such a purpose-specification would be self-defeating since for "developing a self-learning, self-amending algorithm that can serve any purpose" *any and all* data can be regarded as "necessary" and "relevant".

However, as I will note below, at 2.2, the AI Act, rather than banning GPAI models, merely "*aims to ensure that general-purpose AI models are safe and trustworthy*". This even applies to "General-Purpose AI Models with Systemic Risks": they may be highly dangerous and are therefore strictly regulated – but they are not banned.[36] It may require case-law from the Court of Justice to resolve the tensions between the AI Act and the GDPR in this regard.

In fact, as I will note in section 3, the Digital Omnibus proposals only confuse this issue further.

The second issue that arises in this context is whether, or when, "training" of an AI model can be said to be "compatible" with the purpose for which the data are collected, either in general or if such "training" can be said to constitute "scientific research" of the kind covered by Article 89 GDPR.

On the first point, I believe that it will generally *not* be possible to argue that it is necessary to develop an AI model for any specific purpose, say medical treatment (although that may change in future). For the time being, most societal issues still are and can be addressed without the need for AI.

Rather, the most appropriate tests for the time being will be (i) whether, or when, the development and "training" of an AI model can be said to constitute "scientific research", and (ii) if so, whether "appropriate safeguards" (or conditions) have been put in place to protect "the rights and freedoms of the data subject[s]" (as required by Article 89(1)).

On the first point, Kuner *et al.* rightly note that:[37]

> [T]he phrase 'in the public interest' in Article 89 modifies only 'processing for archiving purposes'; thus, data processing for scientific or historical research purposes or statistical purposes is covered by the Article even if they are not 'in the public interest.'

"Scientific" research can therefore include "scientific research" by a private company for commercial purposes. However:[38]

> [T]he derogations [from the main purpose limitation principle set out in Article 5(1)(b)] only apply to the purposes mentioned in the respective paragraphs.[39] This means, for example, that when a company conducts scientific research then such processing may fall under a derogation, but its subsequent use of the research results for commercial purposes would not …

To which they rightly add that:[40]

> … in practice, it may not always be easy to differentiate between the original research purposes of data processing and any subsequent [e.g., commercial] purposes.

Earlier this year, the European Data Protection Board issued guidelines on the issue listing the "*Key-indicative factors for determining whether processing of personal data is motivated by scientific research purposes*", with examples of what, in its view, does or does not constitute "scientific research".[41] Briefly, the "key factors" are:[42]

i. that the research activities are conducted following **a methodical and systematic approach** of the relevant research field;
ii. adherence to **ethical standards** to prevent harm to individuals, ensure respect for human autonomy and consent, transparency, accountability and oversight;
iii. **verifiability** and **transparency**;
iv. **autonomy** and **independence**;
v. that the research aims at "**contributing to the growth of society´s general knowledge and wellbeing**" (although "[t]his does not exclude that the research may also aim to further commercial interests"); and
vi. that the research activities are "**merited** [in that] they have the potential to contribute to existing scientific knowledge or apply existing knowledge in novel ways".

The Board notes in these respects that:[43]

> If the research activities under assessment meet all of the[se] factors …, they can be presumed to constitute scientific research. If the research activities do not meet all factors, the controller needs to justify and be able to demonstrate why the activities should nonetheless be considered scientific research, within the meaning of the GDPR. When doing such an assessment, the more key-indicative factors are present, the more likely it is that the activities constitute scientific research within the meaning of the GDPR.

The guidelines give only one example of AI-related research, in relation to which it says the "key factors" are all met.[44] It relates to a company carrying out research in the field of Generative AI (i.e., not GPAI):[45]

> **Example 2: Processing of personal data for scientific research purposes**
>
> GSPM AI GmbH is a start-up company that conducts research and development in the field of generative AI and is looking to conduct research on the topic of bias in generative AI models. To do so, GSPM AI intends to apply for funding for its research from a foundation which provides funding in various fields of technological research. The conditions for funding include that research projects must use established scientific methods, be subject to ethical review and the results of the research must be made publicly available. The choice of scientific method, formulation of hypotheses etc. is to be determined by the researchers applying for sponsorship. In so far as the intended research fulfils the conditions for funding, the research activities will correspond to factor i (methodical and systematic approach) as well as factor iv (autonomy and independence).
>
> GSPM AI partners with a university faculty that conducts research in the field of robotics and AI and that is looking to get involved in research concerning bias in generative AI models. By doing so, the research project will include qualified researchers, in line with factor iv (autonomy and transparency).
>
> The company and the university set up a joint research group that starts to plan for the research project, applies for sponsorship and submits to review by the university's ethical board, which is composed of independent experts. Subjecting the research project to ethical review corresponds to factor ii (adherence to ethical standards).
>
> The joint research group is granted sponsorship. Having successfully completed the project, GSPM AI and the university faculty publish a peer-reviewed paper, in line with factor iii (verifiability and transparency). By sharing the research results through publication, the research projects benefit the larger research community, while not excluding commercial use of the results by GSPM AI. Therefore, the research activities correspond to factor v (objectives of the research).
>
> In sum, the processing of personal data within the scope of the research project takes place for scientific research purposes, in the meaning of the GDPR.

The following appear to follow from the above:

It would appear to be clear – if only from the fact that Article 89 constitutes an exception to, or at least an important qualification of, the crucially important purpose-limitation principle – that although the *concept* of scientific research should be broadly applied,[46] the

exception/qualification as such should be narrowly applied, to genuine scientific research and research purposes that meet the EDPB's "key factors" only – and not stretched to cover commercial *use* of a trained AI model.

In practice, in my opinion, although this will always rest on a case-by-case assessment, most of the AI model developments and trainings carried out by the leading "Big Tech" companies does <u>not</u> meet the EDPB criteria, and their commercial *uses* of their models is in any case never protected by Article 89 GDPR.

Of course, even when Article 89 can be relied on, there is still the matter of what **safeguards** are "appropriate" in relation to (genuinely scientific) AI research. Notably, the AI Act simply does not apply (at all) to "*AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development*" (Article 2(2)(6)). In this regard, the GDPR is therefore the main point of reference.

However, the GDPR as such leaves this matter to quite a large extent to Member States and indeed to entities involved in the activities mentioned (or their sectoral organisations). Often, Member States regulate matters such as archiving and certain types of scientific research (e.g., medical research) by law. In such cases, the relevant laws, or regulations adopted under those laws, should also set out the relevant safeguards. But the text also allows for the adoption of the relevant rules and safeguards by those carrying out the archiving or scientific research, e.g., in sectoral codes of conduct – although, presumably, these should be binding and enforceable. In that respect, two matters are notable:

First of all, the question whether any safeguards that apply in relation to any specific (genuinely scientific) AI research project or activity are sufficient or "appropriate" is judiciable. National rules on data use for scientific AI model training with insufficient safeguards to protect the data subjects are in breach of the GDPR and can and should be struck down by the courts (and if needs be ultimately by the EU Court of Justice).

Secondly, the "*key indicative factors for determining whether processing of personal data is motivated by [genuinely] scientific research purposes*", adduced by the EDPB, should provide the most important guidance on this matter: if the research parameters do not meet those EDPB standards, the research is not covered by Article 89 and is not subject to appropriate safeguards.

Finally, I should note that with regard to "research" relating to AI including the training of AI models that is *not* "genuinely scientific" as judged by the EDPB standards – such as most AI research and training by the commercial "Big Tech" companies – neither the last part of Article 5(1)(b) nor Article 89 GDPR applies: such research is not "compatible" with the original purpose for which the data are collected. Therefore, all the normal requirements of the GDPR (as discussed in the present section of this paper) should be fully complied with. In particular, as further discussed below, if sensitive data are included in the relevant large data sets (which as I have already noted is virtually inevitable), this means that the processing of personal data for such non-scientific "research" and training cannot be based on the "legitimate interest" legal

basis in Article 6 as such, but rather, must meet one of the special criteria set out in Article 9 GDPR.

iv. Legal bases for the processing of personal data generally, and of sensitive data in particular

(a) General

Under the GDPR, any processing of personal data must be based on one of the **legal bases** listed in Article 6. The main legal bases listed are: that the data subject gave her or his consent; that the data are **necessary** for the performance of a contract (or pre-contractual checks); or to comply with a legal obligation; or to protect a person's vital interests; or to perform a task in the public interest or in the exercise official authority (Article 6(1)(a) – (d)); or finally that:

> [the] processing is necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data. (Article 6(1)(f))[47]

However, Article 9(1) in principle prohibits processing of certain "special categories of personal data" (usually referred to as "**sensitive data**"), subject to limited exceptions set out in Article 9(2). If such data are being collected and/or otherwise processed it does not suffice that one of the Article 6 legal bases can be invoked; rather, *in addition* one of the limited exceptions in Article 9(2) must also apply.

Below, I first discuss the broad application that is given to Article 9(1) in the guidance and the case-law and note that in the development and training of AI models it is in practice all but unavoidable that identifiable ("single-out"-able)[48] sensitive data are also collected. This means that such development and training must in practice (pretty much) always be based on an Article 6 legal basis and fall within one of the situations (exceptions) listed in Article 9(2).

(b) The broad application of Article 9(1) GDPR

Article 9(1) lists as "special categories of personal data": data that reveal or can reveal (or can be used to infer: see below) "racial or ethnic origin,[49] political opinions, religious or philosophical beliefs, or trade union membership", genetic data, biometric data (if used for the purpose of uniquely identifying a natural person), data concerning health and data concerning a natural person's sex life or sexual orientation.

As the EDPB notes, with reference to the CJEU judgment in *Meta v. Bundeskartellamt*:[50]

> **[A] set of data that contains at least one sensitive data item is deemed sensitive data in its entirety**, in particular if it is collected *en bloc* without it being possible to separate the data items from each other at the time of collection.[51] Further, it should be recalled that data are deemed sensitive if such data allow information falling within one of the categories referred to in Article 9(1) GDPR to be revealed.[52] **It is irrelevant whether or not the information revealed by the processing operation in question is correct and whether the controller is acting with the aim of obtaining information that falls within one of the special categories referred to in that provision.[53] Hence, according to the jurisprudence of the CJEU, the relevant question is whether it is objectively possible to infer sensitive information from the data processed, irrespective of any intention of actually doing so.**

This makes clear that the in-principle prohibition of processing of sensitive data must be broadly applied: controllers cannot excuse the processing of such data by claiming that their data sets only contain some minor amounts of sensitive data, (more or less) accidentally or unintentionally collected; or that they did not specifically seek to collect such data and/or did not try and check their accuracy.

More specifically: a controller can also not claim to not be processing sensitive data if they filter out those data immediately after they are collected and before using them, because they are still collected, which is a form of processing – to which therefore the whole of the GDPR applies.[54]

Moreover, the data do not have to explicitly reveal the sensitive matter: it suffices if the relevant sensitive characteristic *can be inferred* from the data. For instance, much can be inferred from a record of location data relating to a person over time. It can show that they regularly attend a certain medical clinic,[55] or a certain mosque, church, synagogue or temple, or political meetings. One can infer from such data that the person in question (probably) suffers from a relevant ailment and/or is seeking some treatment (such as an abortion), or is a member of the relevant religious group, or a supporter of the relevant political organisation (although of course they could be a medical professional, or a cleaner without any ailment or of a different faith, or a [uniformed or undercover] police office or security agent). However, there is little or no guidance on what level of certainty is required in relation to such inferences. The Article 29 Working Party opined that:[56]

> [t]he term "data *revealing* racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership" is to be understood that not only data which by its nature contains sensitive information is covered by this provision, but also data from which sensitive information with regard to an individual can be **concluded**. –

without discussing the level of confidence that would be required to underpin the conclusion (or inference). It also noted that:[57]

> Photos and images of persons, such as those published on the Internet or taken by traffic monitoring or other surveillance cameras, are especially problematic. Since images can reveal information about an individual's ethnic origin or health status, for example, they may be considered sensitive data as defined in Art. 8(1) of [the 1995 Data Protection Directive] [or Article 9(1) of the GDPR – DK], with the result that these data may not be processed without the individual's consent [or unless another condition in Article 9(2) is met – DK].

I am dwelling on this here because it is in practice as good as impossible to collect the kinds of massive data sets that are used to train Large Data AI Models without at least some identifiable ("single-out"-able) sensitive data (as defined above) being "hoovered" up (collected "*en bloc*" to use the Court's terminology).

At the very least, it will be possible to infer sensitive characteristics from such very large data sets: researchers have, for instance, been able to infer a person's sexual orientation from their Spotify play list. Even if no such inferring is intended, and even if those sensitive data only form a small part of the large data set in question, and even if the sensitive data were to be filtered out immediately after collection (if that were indeed possible), this means that in practice Article 9 GDPR will (almost?) always have to be complied with in the creation of very large data sets to develop and train an AI model. In other words:

The creation of very large data sets to develop and train an AI model cannot (almost never?) be based on the "legitimate interest" legal basis by itself: in (almost?) all cases, a condition set out in Article 9(2) must also be met.

(c) Meeting the requirements of Article 9 GDPR in the development and training of AI models and especially in the development of General-Purpose AI models

In my opinion, it can, at least for now, *not* be argued that the creation and training of specific- or general-purpose AI models can ever be said to be "**necessary**" in any of the contexts, or for any of the aims, listed in Article 9(2)(b) – (d), (f), (h) or (i), i.e.: to meet requirements of employment or social welfare legislation; to protect the vital interests of a data subject or another person; to process data by not-for-profit bodies with political, philosophical, religious or trade union aims on their members; for use in legal claims; for purposes of preventive or occupational medicine etc., or for wider public health reasons.

I also believe that, again at least for now, it is difficult to see how it could be argued that the creation and training of specific- or general-purpose AI models can be said to be "necessary" for any other "reasons of substantial public interest" (Article 9(2)(g)).

As noted at *iii.*, above, for the time being, most societal issues still are and can be addressed without the need for AI.

It follows from the Court of Justice's judgment in the case of *Schrems v. Meta (Facebook)* that developers of AI models can also not rely on the "manifestly made public by the data subject"

exception in Article 9(2)(e)) to "hoover up" sensitive personal data from publicly accessible sources such as online newspapers or social media to train their models.[58]

In theory, sensitive data can be used for the development and training of AI models with the consent of the data subject. However, such consent would have to be "explicit" as well as free, informed and specific etc. (Article 9(2)(a) read with the definition of "consent" in Article 4(11); see also Article 7). Moreover, here another issue arises, again especially in relation to collecting and further processing of sensitive data for the development and training of General-Purpose AI Models. This is that, in line with what I said about purpose-specification and -limitation at *iii*., above, in my opinion "consent" for "the training of a general-purpose AI model" cannot be said to be given for a sufficiently specific, "specified" purpose to meet the consent requirements of Article 6(1)(a) as further clarified in Articles 4(11) and 7. Such a purpose-specification would again be "too vague or general", and the consent therefore not sufficiently "specified" to meet the requirements of Article 9(2)(a).

People could consent to the use of their personal data including sensitive data for *genuine* "scientific AI research purposes" in relatively broad terms, e.g., "for the development of an AI model for medical research" – although it would be at least advisable to ask such volunteers for more specific consent once a specific medical research purpose was chosen (e.g., "research into cervical cancer").[59] But as also explained at *iii.*, above, this could only cover genuine research meeting the EDPB "key factors" and not the commercial use of the model.

This also applies to scientific research based on (read: authorised by) Union or Member State law, listed in Article 9(2)(j), because this too only covers *genuine* "scientific research".

Moreover, although as also noted at *iii.*, above, the matter of what constitute "appropriate safeguards" in terms of Article 89 – or "suitable and specific measures to safeguard the fundamental rights and the interests of the data subjects" under Article 9(2)(j) – is, to quite a large extent, left to Member States and national research bodies – but as noted above, the relevant laws and (binding) science body rules should at least require compliance with (as many as possible) "key factors" indicated in the EDPB guidelines.

In section 2.2., below, I will discuss further, with reference to the non-application of the AI Act to:

> AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development. (Article 2(6))

v. Profiling and automated decision-making[60]

(a) *Profiling and AI*

Article 22(1) GDPR stipulates that in principle (subject to exceptions set out in Article 22(2) and strict conditions – and an important further prohibition – set out in Article 22(3) and (4)):

> The data subject shall have **the right not to be subject to a decision based solely on automated processing, including profiling**, which produces legal effects concerning him or her or similarly significantly affects him or her.

This makes clear that any decision based solely on profiling constitutes a "*decision based solely on automated processing*". Article 4(4) GDPR defines the concept as follows:

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

This – profiling – is of course precisely what AI models and more especially Large Data AI Models including Large Language Models do: they find correlations in large data sets that are linked to specific "aspects" of persons – e.g., that persons who used two different medications that were separately prescribed often suffered from certain complications; use this to conclude that there is a high probability that future persons who are prescribed those medications will also develop those complications; and therefore recommend that those two medications should not be simultaneously prescribed.

Often, such models will produce "**scores**" that indicate the overall significance of the factors that were taken into account (and weighted in various ways) in relation to a matter on which a decision is to be taken, such as whether or not to offer a job to someone, or admit someone to a school or university course, or to release a prisoner on parole.[61] In self-learning AI, the algorithm that leads to the score is furthermore constantly adjusted by the AI itself, i.e., the AI and the score are "***dynamic***". It will often be difficult if not impossible for even the developer of the model to fully understand how the (dynamic) score is arrived at in a particular case; or the developer of the model may refuse to disclose the precise logic used on the ground that it is a "commercial-" or "trade secret". In that case, the model is in effect a "**black box**", the output of which is effectively ***unreviewable*** and ***unchallengeable***, at least for the entity that relies on the score and for the data- and decision subject.

(b) *"Black box" systems*

As already mentioned at *ii.*, above, and discussed in more detail in my earlier paper on the important "transparency", "explainability" and "procedural fairness" requirements that flow from Article 22, **in its *PNR* judgment, the Court effectively prohibited the use of certain self- or machine-learning (i.e., self-modifying) AI systems ("ML/AI systems") – what an UK Information Commissioner's Office/Alan Turing Institute paper referenced in that previous paper calls "black box" systems – in the taking of decisions under the PNR directive, in terms that also apply in relation to Article 22 GDPR**.[62]

(c) *Scoring*

Furthermore, in its *SCHUFA* judgment, the Court of Justice held that if information relating to an individual is used to create such a "**score**" (in that case a credit score), and this score is then decisively relied on by a decision-taker, that score-creation will already constitute the taking of an automated decision in terms of the GDPR, even if the entity that calculated the score is different from the entity that relies on the score in its decision-taking.[63] It is worth noting the outline of the case as provided in that judgment:

> SCHUFA is a private company under German law which provides its contractual partners with information on the creditworthiness of third parties, in particular, consumers. To that end, it establishes a prognosis on the probability of a future behaviour of a person ('score'), such as the repayment of a loan, based on certain characteristics of that person, on the basis of mathematical and statistical procedures. The establishment of scores ('scoring') is based on the assumption that, by assigning a person to a group of other persons with comparable characteristics who have behaved in a certain way, similar behaviour can be predicted.
>
> It is apparent from the request for a preliminary ruling that OQ was refused the granting of a loan by a third party after having been the subject of negative information established by SCHUFA and transmitted to that third party. OQ applied for SCHUFA to send her information on the personal data registered and to erase some of the data which was allegedly incorrect.
>
> In response to that request, SCHUFA informed OQ of her score and outlined, in broad terms, the methods for calculating the scores. However, referring to trade secrecy, it refused to disclose the various elements taken into account for the purposes of that calculation and their weighting. Lastly, SCHUFA stated that it limited itself to sending information to its contractual partners and it was those contractual partners which made the actual contractual decisions.
>
> (Paras. 14 – 16)

The Court's response was clear:

> Article 22(1) of the GDPR must be interpreted as meaning that the automated establishment, by a credit information agency, of a probability value based on personal data relating to a person and concerning his or her ability to meet payment commitments in the future [already in itself] constitutes 'automated individual decision-making' within the meaning of that provision, where a third party, to which that probability value is transmitted, draws strongly on that probability value to establish, implement or terminate a contractual relationship with that person.
> (Para. 73, repeated verbatim in the conclusion in para. 75. Words in square brackets added)

This means that:

AI models that create scores that are relied on by other parties to take "significant" decisions on individuals are in principle prohibited and only permitted under the exceptions and conditions set out in Article 22(2) and (3) and the prohibition in Article 22(4).

(d) *When automated decision-making can be allowed:*

Article 22(2) allows the taking of fully-automated decisions – and therefore any decisive score-creation (provided it is not done in a "black box") – in only three cases, i.e., if the processing (or the score-creation and score-reliance):

(a) is ***necessary*** for entering into, or performance of, a **contract** between the data subject and a data controller;

(b) is authorised by Union or Member State **law** to which the controller is subject and which also lays down ***suitable measures to safeguard the data subject's rights and freedoms and legitimate interests***; or

(c) is based on the data subject's explicit **consent**.

As to Article 22(2)(a), I cannot see how it will ever be "***necessary***" for an entity to take automated decisions (including score-based decisions) in order to decide on whether to enter into a contract with a data subject. A company may well demand a prospective client's agreement to such decision-making, but that does not make it objectively "necessary". I can imagine circumstances in which the performance of a contract may involve some automated actions, e.g., when games software "decides" to move a player to a higher level. But this will have to be applied very strictly.

I will come to the other two situations below.

(e) *What safeguards should be put in place?*

Article 22(3) stipulates that:

> In the cases referred to in points (a) and (c) [of Article 22(1), quoted above], the data controller shall implement **suitable measures to safeguard the data subject's rights and freedoms and legitimate interests**, [including] at least ***the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision***.

This requires **transparency** and **an effective remedy** involving **human intervention** (a "***human in the loop***"). Moreover, any law that authorises automated decision-making (Article 22(b)) can also not be said to include "suitable measures to safeguard the data subject's rights and freedoms and legitimate interests" if it did not also provide these matters.

(ea) Transparency

One crucial precondition for the exercise of the rights of a decision subject under Article 22(3) to "express his or her point of view" and to "contest" an automated decision is compliance with the **transparency** requirements of Article 13(2)(f) and 14(2)(g), already briefly discussed

above, at *ii.*, above, i.e., that the data (and decision-) subject is informed about ***the fact that automated decision-making is being planned***; about the "***logic***" that is to be used in the automated decision; and about "*the envisaged **consequences** of such processing for the data subject*".

Moreover, since those matters are also an inherent part of the **right to an effective remedy** in relation to any matter that affects a person's rights and freedoms, guaranteed by Article 47 CFR, those same transparency/information requirements must also apply to automated decision-making authorised by law under Article 22(2)(b).

In sum:

To the extent that automated decision-making is allowed under Article 22(2) (as exceptions to the in-principle prohibition in Article 22(1)), i.e., in a contractual context, with the data subject's consent, or when authorised by law, the data subject must still always be fully informed of all the matters listed in Article 13(2)(f) and 14(2)(g) including in particular of the "logic" used in the decision.

Crucially, they must be informed of this "logic" also if the data underpinning this "logic" are not "personal data" within the meaning of Article 4(1) GDPR and Article 8 CFR. This is the *caveat* I mentioned at the beginning of this section. I will return to that at 3.2, below.

(eb) The level of detail to be provided

In its *SCHUFA* judgment, the Court did not address the question of **the level of detail that is to be provided** about how "scores" are arrived at in order for that information to be sufficiently "meaningful" to meet the Articles 13 and 14 transparency- and Article 47 effective remedy requirements. However, the Advocate General did address this in his opinion on the case in terms that I believe still represents the correct interpretation and application of the law. Specifically, as I concluded in the previous paper, already mentioned:[64]

In his opinion in *SCHUFA*, the Attorney General recognises the direct link between the right under Articles 13, 14 and 15 GDPR to "*meaningful information about the logic*" used in the taking of automated decisions and the right under Article 22(3) GDPR to "contest" the decision, which is itself directly linked to the right to an effective remedy under Article 47 CFR.

In particular, <u>the information to be provided to the data subjects under Articles 13, 14 and 15 GDPR, but in particular in response to a subject access request under the latter article, should be sufficient to allow a meaningful review of whether the outputs of the AI system were fair</u> – which, as I have shown [in sub-section *ii.*, above], means that those outcomes

should not be biased against protected classes of people or produce excessive false negatives or false positives (FN/FP rates).

Consequently, at least in response to a subject access request, the data subject should be provided with sufficient information on the factual elements taken into account by the AI system and the weighting of those elements, and with sufficient information on any such biases and FN/FP rates to allow for a meaningful review, first by any "human in the loop" within the organisation of the controller (taking into account the views of the data subject) (see below, at (*ec*)), and second, if the matter is not resolved at that level, by any tribunal that deals with the case.

Failure to provide adequate information on an AI system, and/or on the outputs of an AI system, to allow for such reviews would, as the Attorney General notes, compromise the essence of the fundamental right to data protection and also the essence of the fundamental right to an effective remedy.

The above is in line with another German case (also mentioned in that paper), in which a person (who by the way had a good SCHUFA score) was denied a credit card by a bank on the basis of an automated score, based on criteria set by the bank, but with the bank refusing to provide details of the scoring algorithm. In that case, the Berlin data protection authority held that:[65]

> When companies take automated decisions, they are obliged to explain these in a sensible and verifiable way. The affected person must be able to verify the automated decision. … In relation to an automated decision on an application for a credit card, a bank is obliged to inform its customers about the underlying reasons for a refusal. This includes [i.e., the bank must provide the customer with] concrete information about the data basis and the decision elements as well as the criteria for the refusal in the specific case. (My translation)

I should add that, in my opinion, if a person is not provided with the "sufficient information" outlined by the Advocate General and the Berlin data protection authority, that person can also not give his or her "*informed* consent". Obtaining "consent" to processing – and *a fortiori* to automated decision-making – from a person who cannot understand what exactly they are consenting to, is not valid consent in terms of the GDPR.[66]

(ec) "Human intervention"

Article 22(3) makes clear that in relation to automated decisions taken in a contractual context or with the consent of the data subject, the data (and decision-) subject has the right "to obtain human intervention" in the context of them "contesting" the decision. In my opinion, since this issue, too, is intimately connected to the right to a remedy under the Charter, this must also apply to automated decision-making authorised by law.

In a postscript to my paper on algorithmic discrimination, I discussed a good – i.e., bad – example of proposed UK rules that do not meet this condition, noted by Graham Smith.[67] This concerns proposed guidance to U2U platforms and search engines on "*proactive technology: automated systems intended to detect particular kinds of illegal content and content harmful to children, with a view to blocking or swiftly removing them*", under the UK Online Safety Act, issued for consultation by the UK communications services regulator, Ofcom.[68]

Apart from noting that the proposed rules do not meet the European Court of Human Right's "quality of law" requirements (that legal rules must be clear, precise and foreseeable in their application), he also noted that the rules envisaged "prior constraint" on free speech; that they accepted that some forms of legitimate speech would be wrongly suppressed; and – crucially in the present context – that although the scheme mandated a first and second review of contested decisions, both could be done by the system *without human involvement*.

It is of course preposterous to believe that AI can make the very difficult, context-specific assessment required to judge whether some expression is or is not lawful. But this scheme makes it worse by allowing suppression of free speech by fully-automated means without any human assessment.

This example shows the importance of the Article 22(3) safeguards, and of the "human in the loop" requirement in particular.

(ed) Automated decisions based on sensitive data

Finally, Article 22(4) stipulates that:

> Decisions referred to in paragraph 2 shall not be based on special categories of personal data referred to in Article 9(1) [i.e., sensitive data], unless point (a) or (g) of Article 9(2) applies and suitable measures to safeguard the data subject's rights and freedoms and legitimate interests are in place.

Points (a) and (g) of Article 9(2) relate respectively to processing (i.e., here: the taking of automated decisions) on the basis of the "**explicit consent**" of the data subject, and to processing (the taking of automated decisions) that is:

> necessary for reasons of **substantial public interest**, on the basis of **Union or Member State law** which shall be ***proportionate*** to the aim pursued, ***respect the essence of the right to data protection*** and provide for ***suitable and specific measures to safeguard the fundamental rights and the interests of the data subject***.

In this respect it is crucial to recall **the absolute prohibition of discrimination**, discussed at *ii*., above. Article 22(4) may of course not be read as an exception to that universally fundamental rule:

> **Individuals cannot validly consent to being subjected to automated decision-making that discriminates against them, and states may not adopt laws that allow for the taking of automated decisions that (even inadvertently) lead to discrimination.**

In that regard, it is important to recall that international law defines discrimination as:[69]

> *"**Any distinction, exclusion, restriction or preference** based on [any ground such as sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation] **which has the purpose or effect of nullifying or impairing the recognition, enjoyment or exercise, on an equal footing, of human rights and fundamental freedoms in the political, economic, social, cultural, civil or any other field of public life.**"* (emphases added)

It follows from this that:

> **To the extent that any EU or Member State laws do or in future will allow the use of sensitive data in automated decision-making, the Union or the Member States in question must ensure that any biases in the *output* of the relevant AI systems are identified and any biases in the consequent *outcomes* fully eliminated.**[70]

vi. Anonymising personal data is a form of processing of such data

> **Anonymising personal data is a form of processing of such data because it involves "adaptation or alteration" of the original (still identifiable) data .[71] This means that the action of anonymising data must meet all and any of the requirements of the GDPR. In particular, the anonymising must serve a specific, specified purpose; it must have a legal basis in Article 6 and, if the pre-anonymisation data are sensitive, it must meet one of the situations or conditions of Article 9; and the disclosure of the supposedly anonymised* data to a third party still involves "processing of [identifiable] personal data" by the entity that discloses the data (because that entity can still link the supposedly anonymised* data to specific [named or otherwise singled-out] individuals, even if the recipient of the supposedly anonymised data may not be able to do this).**
>
> *** I am saying "supposedly anonymised" here because in practice it is becoming very difficult to really fully anonymise personal data.**

This issue is currently being addressed in a consultation on guidelines on anonymisation that is being held by the European Data Protection Board. I will discuss these in section 3.1, below, with further comments.

vii. The obligation to carry out a Data Protection Impact Assessment (DPIA)

Finally, I should note that Article 35(1) of the GDPR stipulates that:

> Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. ...

Article 35(3) adds that:

> A data protection impact assessment referred to in paragraph 1 shall in particular be required in the case of:
>
> (a) a systematic and extensive evaluation of personal aspects relating to natural persons which is based on automated processing, including profiling, and on which decisions are based that produce legal effects concerning the natural person or similarly significantly affect the natural person;
>
> (b) processing on a large scale of special categories of data referred to in Article 9(1), or of personal data relating to criminal convictions and offences referred to in Article 10; or
>
> (c) a systematic monitoring of a publicly accessible area on a large scale.

This is not the place to discuss DPIAs in detail. Suffice it to note that the Article 29 Working Party adopted guidelines on this in the run-up to the GDPR that were endorsed by the European Data Protection Board once the GDPR came into force.[72] These note that the GDPR sets out the minimum features of a DPIA in Article 35(7) and recitals 84 and 90. They must contain, at least:

- a description of the envisaged processing operations and the purposes of the processing;
- an assessment of the necessity and proportionality of the processing;
- an assessment of the risks to the rights and freedoms of data subjects; and
- the measures envisaged to address the risks and demonstrate compliance with the GDPR.

I am mentioning this here because there is a requirement for a different (if somewhat overlapping) Fundamental Rights Impact Assessment in the AI Act. I discuss this further in the next section.

Overall conclusion on the GDPR "guardrails" for AI-based automated decision-making:

The above analyses make clear that the GDPR as interpreted by the Court of Justice and further clarified in the guidance issued over the years by the Article 29 Working Party and the European Data Protection Board imposes strict limits and conditions on AI-based automated decision-making. In the next section, I discuss how the AI Act and the Digital Omnibus proposals compare and relate to – or weaken – those GDPR "guardrails".

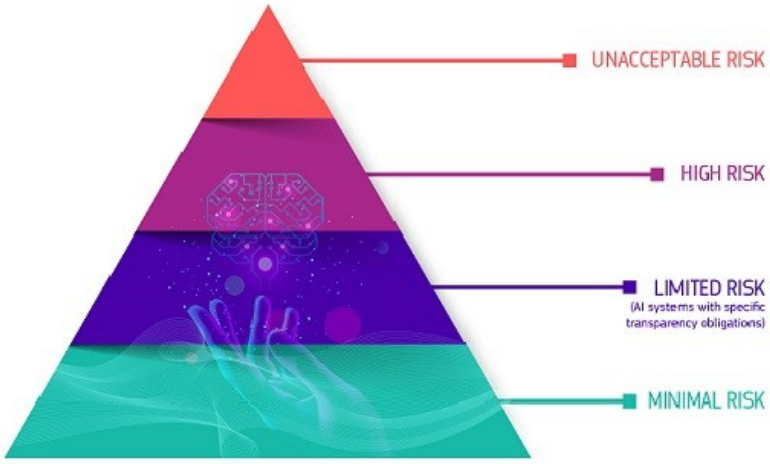
2.2 The EU AI Act

i. Introduction and overview

The EU AI Act (Regulation (EU) 2024/1689) was published in the Official Journal (OJ) of the European Union on 12 July 2024.[73] It formally "entered into force" on 1 August of that year, but is being brought into actual application in phases: its general provisions (definitions and provisions on AI literacy) came into application on 2 February 2025; the rules for general-purpose AI (GPAI) models on 2 August 2025; and the majority of its other provisions (including those on transparency) on 2 August 2026, when enforcement concerning GPAI models, prohibitions, transparency rules and AI literacy also started. Some new prohibitions for AI systems generating non-consensual sexual deepfakes and child sexual abuse material will apply from 2 December 2026, while the rules for high-risk AI systems in Annex III and the rules for high-risk AI embedded in regulated products covered by Annex I will only come into application in December 2027 and August 2028, respectively.[74]

The European Commission has produced a page that summarises the rules and requirements of the Act as well as a more elaborate "*AI Act Single Information Platform*" with further links, to which I refer the reader for details.[75] [76] Here, the following broad overview (largely taken from those pages) will suffice before I turn to some specific issues most relevant to the present paper.

The Act adopts a "**risk-based approach**" by defining four levels of risk for AI systems,[77] with corresponding descending requirements, as follows:[78]



Unacceptable risk

All AI systems considered a clear threat to the safety, livelihoods and rights of people are banned. **The AI Act prohibits eight practices**, namely:[79]

1. harmful AI-based manipulation and deception
2. harmful AI-based exploitation of vulnerabilities
3. social scoring
4. Individual criminal offence risk assessment or prediction
5. untargeted scraping of the internet or CCTV material to create or expand facial recognition databases
6. emotion recognition in workplaces and education institutions
7. biometric categorisation to deduce certain protected characteristics
8. real-time remote biometric identification[80] for law enforcement purposes in publicly accessible spaces

The prohibitions became effective in February 2025.

But note that the Act does not prohibit either unreviewable and unchallengeable "black box" AI systems (even though the *PNR* judgment makes clear that these are incompatible with the rights to data protection and to an effective remedy) or General Purpose AI systems (even though, at least in my opinion, these are inherently incompatible with the purpose-specification and -limitation principle and thus also with Article 8 of the Charter.)

High risk

AI use cases that can pose serious risks to health, safety or fundamental rights are classified as high-risk. These **high-risk** use-cases include:

- AI safety components in critical infrastructures (e.g. transport), the failure of which could put the life and health of citizens at risk
- AI solutions used in education institutions, that may determine the access to education and course of someone's professional life (e.g. scoring of exams)

- AI-based safety components of products (e.g. AI application in robot-assisted surgery)
- AI tools for employment, management of workers and access to self-employment (e.g. CV-sorting software for recruitment)
- Certain AI use-cases utilised to give access to essential private and public services (e.g. credit scoring denying citizens opportunity to obtain a loan)
- AI systems used for remote biometric identification, emotion recognition and biometric categorisation (e.g. AI system to retroactively identify a shoplifter)
- AI use-cases in law enforcement that may interfere with people's fundamental rights (e.g. evaluation of the reliability of evidence)
- AI use-cases in migration, asylum and border control management (e.g. automated examination of visa applications)
- AI solutions used in the administration of justice and democratic processes (e.g. AI solutions to prepare court rulings)

High-risk AI systems are subject to **strict obligations** before they can be put on the market:

- adequate risk assessment and mitigation systems
- high-quality of the datasets feeding the system to minimise risks of discriminatory outcomes
- logging of activity to ensure traceability of results
- detailed documentation providing all information necessary on the system and its purpose for authorities to assess its compliance
- clear and adequate information to the deployer
- appropriate human oversight measures
- high level of robustness, cybersecurity and accuracy

The rules for high-risk AI [mainly came] into effect in August 2026; the remaining rules will come into effect in] August 2027.

The *Future of Life Institute* adds to this in another useful high-level summary:[81]

> Most of the text addresses **high-risk AI systems, which are [strictly] regulated [as per the above]** and the majority of obligations fall on providers (developers) of high-risk AI systems.
>
> A smaller section handles **limited risk AI systems, [which are only] subject to lighter transparency obligations**: providers and deployers must ensure that end-users are aware that they are interacting with AI (chatbots and deepfakes).
>
> **Minimal risk is unregulated.**

As the Commission explains in relation to the last category:[82]

> The AI Act does not introduce rules for AI that is deemed minimal or no risk. The vast majority of AI systems currently used in the EU fall into this category. This includes applications such as AI-enabled video games or spam filters.

A special chapter (Chapter V) deals with **General-Purpose AI models**,[83] which (as already noted and critically discussed in relation to the GDPR "purpose specification and -limitation" principle in section 2.1, sub-section *iii.*, above) are defined as follows:

> '**general-purpose AI model**' means an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, *except AI models that are used for research, development or prototyping activities before they are placed on the market* (Article 2(63))

I will discuss the qualification (above in italics) in sub-section *iii.*, below.

There is also a special provision (Article 50) (that actually forms its own chapter, Chapter IV) on **transparency obligations** – but as will note at *iv*, below, these fall far short of the information requirements under the GDPR.

Article 2 makes clear that (like the GDPR**) the AI Act also applies extraterritorially** to providers (developers) of AI systems or General-Purpose AI models who are established or situated outside the EU/EEA if they place those systems on the market in the EU/EEA and to providers and deployers (users) of AI systems where the output produced by the system is used in the EU/EEA (Article 2(1)(a) and (c), as well as to importers and distributors of such systems (Article 2(1)(d)) and "affected persons" (who the GDPR calls "data subjects") who are located in the EU/EEA.

In the next sections, I will discuss in turn:

- the (in relation to data protection, in my view, inappropriate) nature of the AI Act, its AI market- (rather than fundamental rights-) promoting purpose and the rather special (and limited) risks it addresses;
- the scope and especially the sweeping (and in my view, manifestly excessive) exceptions and exemptions set out in the Act, and the implications in relation to the application of the GDPR (to which the Act pays limited service); and
- the (in my view deficient) rules on transparency relating to AI.

ii. The nature and purpose of the AI Act and the risks it addresses

Article 1(1) of the AI Act sets out the subject matter of the Act as follows:

Article 1

Subject matter

> The purpose of this Regulation is to improve the functioning of the internal market and promote the uptake of human-centric and trustworthy artificial intelligence (AI), while ensuring a high level of protection of health, safety, fundamental rights enshrined in the Charter, including democracy, the rule of law and environmental protection, against the harmful effects of AI systems in the Union and supporting innovation.

This makes clear that the AI Act is first and foremost an internal market (and market-promoting) measure, with fundamental rights merely seen as a necessary (and in the view of some, regrettable) limitation.[84] Notably, although the article refers in broad terms to "fundamental rights enshrined in the [EU Charter of Fundamental Rights]" and the rule of law, it does not specifically mention data protection or the right to an effective remedy, in spite of these being amongst the most important fundamental rights affected by AI. "Data subjects" – the entities at the heart of data protection – are only mentioned twice, in Article 59, which is the article that I believe to be, with Article 10(5), the most dangerous and data protection-undermining of all the provisions in the Act (as discussed below, at *iii*).

Moreover, although Article 2(7) stipulates that:

> Union law on the protection of personal data, privacy and the confidentiality of communications applies to personal data processed in connection with the rights and obligations laid down in this Regulation. This Regulation shall not affect Regulation (EU) 2016/679 [the GDPR] or (EU) 2018/1725 [the regulation setting the data protection rules for EU institutions], or Directive 2002/58/EC [the e-Privacy Directive] or (EU) 2016/680 [the Law Enforcement Directive] … –

it adds to this that the above is:

> … without prejudice to Article 10(5) and Article 59 of this Regulation. –

i.e., precisely those dangerous provisions.

In fact, unlike the GDPR (and the other EU data protection laws) which are squarely aimed at protecting fundamental rights,[85] the AI Act is more like a product safety law. This is clear from the definition of "**risk**" by means of a cross-reference in Article 79(1) to the regulation on market surveillance and compliance of products:

> **AI systems presenting a risk** shall be understood as a '***product presenting a risk***' as defined in Article 3, point 19 of Regulation (EU) 2019/1020, in so far as they present risks to the health or safety, or to fundamental rights, of persons.

The regulation referred to, Regulation 2019/1020, stipulates in Article 3(19) that:[86]

> ‘**product presenting a risk**’ means a product having the potential to affect adversely health and safety of persons in general, health and safety in the workplace, protection of consumers, the environment, public security and other public interests, protected by the applicable Union harmonisation legislation, to a degree which goes beyond that considered reasonable and acceptable in relation to its intended purpose or under the normal or reasonably foreseeable conditions of use of the product concerned, including the duration of use and, where applicable, its putting into service, installation and maintenance requirements.

These risks are repeatedly referenced in the AI Act:

- the logs that must be kept in relation to high-risk AI systems must “enable the recording of events relevant for … **identifying situations that may result in the high-risk AI system presenting a risk within the meaning of Article 79(1)** … (Article 12(2)(a));
- **where a high-risk AI system “presents a risk within the meaning of Article 79(1)”**, providers must inform the market surveillance authorities (Article 20(2)); importers must inform the provider, the authorised representative and the market surveillance authorities (Article 23(2)); distributors must inform the provider or the importer of the system (as applicable) (Article 24(2)) and the market surveillance authorities if no corrective action is taken (Article 24(4)); and deployers must inform the provider or distributor and the market surveillance authorities (Article 26(5)).

Crucially, these risks only need to be acted upon if “the product” – i.e., here: the high-risk AI system in question – “**adversely affects**” the issues listed: health and safety, protection of consumers, the environment, public security and other public interests protected by EU harmonisation legislation, to which “fundamental rights” have been added incongruously, almost as an afterthought, “***to a degree which goes beyond that considered reasonable and acceptable in relation to its intended purpose or under the normal or reasonably foreseeable conditions of use of the product concerned***”.

The latter qualification may be suited to the placing of consumer goods such as washing machines or cars on the market – but it is not appropriate to matters affecting fundamental rights. Those rights may only be limited where the limitation serves a “legitimate interest in a democratic society” and only to the extent that the limitation is set out in a clear, precise and in its application foreseeable law and “necessary” and “proportionate” to the aim in question.[87] That is a fundamentally different approach to the one more understandably adopted in relation to health and safety risks of “products” and that allows for “risks” posed by products that are “reasonable and acceptable” in the context of the normal use of products.

This health and safety approach permeates the AI Act. Thus, while as noted above the Act imposes important reporting duties on providers, importers, distributors and deployers of high risk AI systems, these are all aimed at alerting those in the supply chain and the "market surveillance authorities"; there is no requirement that these actors must inform individuals affected by an "unacceptable" risk of their systems. That is presumably left to those authorities (just like they would issue health and safety warnings about unsafe products such as washing machines).

As I will note in the next sub-section, as a consequence various provisions in the AI Act are clearly at odds with the GDPR, the Charter of Fundamental Rights and the EU Treaties – in fact, even where the Act pays lip-service to the GDPR, it still inserts loopholes that are incompatible with it and, in my view, with the Charter and the Treaties.

iii. Scope and exclusions; relation to the GDPR and its sister instruments

(a) *Scope and exclusions*

The Commission says the AI Act is "**the first-ever <u>comprehensive</u> legal framework on AI worldwide**,"[88] but that is simply not true: although it has wide application to AI developers, distributors, importers and deployers including developers and deployers outside the EU/EEA, as noted above, at *i*, it also contains sweeping exemptions and exclusions, set out in Article 2(3), (4), (6) and (8), and one that specifically relates to the GDPR in Article 2(7).

I have addressed the first three, relating to national security, military and defence contexts, and cooperation with non-EU countries (so-called "third states") and international organisations, in some detail in another opinion written at the request of the European Centre for Not-for-profit Law, ECNL ("my ECNL Opinion").[89] Here, it will suffice to briefly summarise part of my analyses and reproduce my conclusions in these respects before turning to the – for the present paper, more immediately relevant – qualified exceptions.

- National security

The first sub-clause in Article 2(3) stipulates that:

> This Regulation does not apply to areas outside the scope of Union law, and shall not, in any event, affect the competences of the Member States concerning **national security**, *regardless of the type of entity entrusted by the Member States with carrying out tasks in relation to those competences*.

The reference at the end to "*entities entrusted by … [EU] Member States with carrying out [national security-related] tasks*" is particularly worrying given the increased "outsourcing" of such tasks to surveillance and data mining companies such as *Pegasus* and *Palentir*.[90]

That aside, in this regard I concluded in my ECNL Opinion that:[91]

> **It follows from the PNR judgment of the CJEU that even when Member States hold exclusive competence in some areas (such as, in particular, national security), if the exercise of that competence affects an area where the EU has competence and that is covered by EU law (e.g., data protection, internal market), the exercise of that exclusive Member State competence may not impinge on the EU legal order or undermine the relevant EU legal rules [such as, in the present context, the rules in the AI Act].**

The exemption in Article 2(3) of the AI Act should be read and applied with the above in mind. It would have been helpful if a recital had expressly clarified this.

Furthermore, although the activities of EU Member States are largely outside the scope of EU law (subject to that qualification), there is no reason why the development and placing on the market of technologies that are to be used for national security purposes should not be subject to EU internal market law. I discuss this under the next heading in relation to military AI, but the same reasoning also applies to national security-related AI.

- Military AI

The remaining clauses in Article 2(3) add further sweeping exemptions, as follows:

> This Regulation does not apply to AI systems where and in so far they are placed on the market, put into service, or used with or without modification exclusively for **military, defence or national security purposes**, *regardless of the type of entity carrying out those activities.*
>
> This Regulation does not apply to AI systems which are not placed on the market or put into service in the Union, where the output is used in the Union exclusively for **military, defence or national security purposes**, *regardless of the type of entity carrying out those activities.*

I note again the reference to "*entities [other than a state's own army or intelligence services] carrying out [military or national security] activities*." This again reflects the increasing normalisation of the outsourcing of such matters to private entities, especially Israeli- or US-based ones with close ties to their rightwing militaristic (and in the case of Israel genocidal) governments.

Recial 24 says that:

> **[This] exclusion is justified both by Article 4(2) TEU and by the specificities of the Member States' and the common Union defence policy covered by Chapter 2 of Title V TEU that are subject to public international law, which is therefore the more appropriate legal framework for the regulation of AI systems in the context of the use of lethal force and other AI systems in the context of military and defence activities.**

To the extent that this sentence suggests that military and defence matters are covered by Article 4(2) TEU and (like national security matters) completely outside the scope of EU law, or could be read like that, this is simply false.

The reference in the remainder of the sentence to Chapter 2 of Title V TEU and to "*public international law [being] the more appropriate legal framework … in the context of the **use** of lethal force and other AI systems in the context of military and defence activities*" is also misleading. Yes, the activities of EU bodies and EU Member States under the European Common Security and Defence Policy (CSDP) (which is part of the EU Common Foreign and

Security Policy, CFSP) and the **use** of force and lethal technology must of course comply with the UN Charter, international humanitarian law and international human rights law. But:

> **This does not mean that the development and placing on the market of technologies to be used in such activities should be outside the scope of internal market rules and more especially internal market rules that also seek to protect fundamental rights guaranteed by the Charter and the Treaties – such as the GDPR (or its sister instruments) and the AI Act.**

The EU **European Defence Action Plan (EDAP)** in fact specifically aims to "***Strengthen the Single Market for defence***"[92] – and trade in military equipment *is* actually already subject to EC directives, most notably the Defence Procurement Directive and the Intra-Community Transfers (ICT) Directive that make up the EU Defence Package (even though they do not appear to be very effectively implemented).[93]

Furthermore, as I also wrote in my ECNL Opinion:[94]

> [B]ecause CSFP activities – like all EU activities – are subject to the Charter, the separate CFSP rules should also unambiguously stipulate that AI systems posing "unacceptable" risks to fundamental rights should not be used by any EU entity, also not "where that use falls under the exclusive remit of the Common Foreign and Security Policy regulated under Title V of the Treaty on the European Union (TEU)." This is all the more important because **unexplainable and unchallengeable AI systems pose an especially serious – indeed, potentially lethal – risk in military contexts.** For instance, they could be used to try and determine whether a certain person in a conflict zone is a combatant (and thus a legitimate target) or a civilian or otherwise protected person; or whether a certain object is a civilian object or a military objective (the making of such distinctions being fundamental to the application of international humanitarian law, IHL).[95]
>
> In fact, there are warnings that ML/AI-based lethal autonomous drones could soon become a reality (if they are not already).[96] As long ago as 2017, the US Department of Defense opened a call for the development of "automatic target recognition of personnel and vehicles from an unmanned aerial system using learning algorithms."[97] The Stop Killing Robots Coalition campaigns against them.[98]

Of course, since I wrote this, we have seen AI being used in targeting people for assassination in Israel's genocide in Gaza in particular:[99]

> **'Lavender': The AI machine directing Israel's bombing spree in Gaza**
>
> The Israeli army has marked tens of thousands of Gazans as suspects for assassination, using an AI targeting system with little human oversight and a permissive policy for casualties.

To leave the development and placing on the market of such disgusting AI systems out of the main EU instrument on AI, the AI Act, is simply inexcusable.

As I concluded in my ECNL Opinion:[100]

> **There is no justification for the exemption in the AI Act relating to AI systems developed [and placed on the internal market] for military [or national security] purposes at all. They should be fully subject to the Act – and to the Charter as interpreted by the Court of Justice.**
>
> **The development and placing on the internal market of military (or dual purposes) goods and services incorporating or using unexplainable and unchallengeable AI systems should consequently be prohibited.**
>
> **Furthermore, the development and placing on the internal market of other military (or dual-use) AI systems should be subject to the same strict conditions as are provided for "high risk" systems in the AI Act, i.e., an ex ante conformity assessment, the establishment of a risk management and data and data governance system, the drawing up and having available for inspection of technical documentation and detailed records, post-market surveillance, etc.**

- AI systems used by third countries or international organisations in the EU Area of Freedom, Security and Justice (AFSJ)

Article 2(4) stipulates that:

> This Regulation applies neither to public authorities in a third country nor to international organisations falling within the scope of this Regulation pursuant to paragraph 1, where those authorities or organisations use AI systems in the framework of **international cooperation or agreements for law enforcement and judicial cooperation** with the Union or with one or more Member States, *provided that such a third country or international organisation provides adequate safeguards with respect to the protection of fundamental rights and freedoms of individuals.*
>
> (The words in italics were added to the Commission draft in the subsequent legislative process. Their omission from the original text is again telling.)

The justification is provided in recital 22. This says that, although (as explained earlier) the AI Act applies also to "providers and deployers of AI systems that are established in a third country, to the extent the output produced by those systems is intended to be used in the Union":

> Nonetheless, to take into account existing arrangements and special needs for future cooperation with foreign partners with whom information and evidence is exchanged, this Regulation should not apply to public authorities of a third country and international organisations when acting in the framework of cooperation or international agreements concluded at Union or national level for law enforcement and judicial cooperation with the Union or the Member States, provided that the relevant third country or international organisation provides adequate safeguards with respect to the protection of fundamental rights and freedoms of individuals. Where relevant, this may cover activities of entities

> entrusted by the third countries to carry out specific tasks in support of such law enforcement and judicial cooperation.

Note yet again the reference to "entities entrusted" – here: by third countries! – to carry out law enforcement tasks. Think again *Pegasus* and *Palentir* and their ilk.

Recital 22 continues to note that:

> Such framework for cooperation or agreements have been established bilaterally between Member States and third countries or between the European Union, Europol and other Union agencies and third countries and international organisations –

and recommends that:

> The authorities competent for supervision of the law enforcement and judicial authorities under this Regulation should assess whether those frameworks for cooperation or international agreements include adequate safeguards with respect to the protection of fundamental rights and freedoms of individuals.
>
> Recipient national authorities and Union institutions, bodies, offices and agencies making use of such outputs in the Union remain accountable to ensure their use complies with Union law.
>
> When those international agreements are revised or new ones are concluded in the future, the contracting parties should make utmost efforts to align those agreements with the requirements of this Regulation.
>
> (Paragraph breaks added for clarity)

In this regard, I concluded in my ECNL Opinion that:[101]

> **[T]he European Data Protection Supervisor should urgently issue an opinion on the use of AI systems in processing of personal data subject to EU data protection rules in relation to which he is competent (Europol Regulation, Law Enforcement Directive, Regulation 2018/1725, etc.), in which he clarifies that [irrespective of the sweeping exemptions in Article 2 of the AI Act] the use of unexplainable and unchallengeable AI systems in the taking of any decision that significantly affects the fundamental rights of individuals is incompatible with the Charter and therefore also with each and every one of those instruments [; and that this also applies in relation to any relevant international agreements].**

More generally, the revision of the extensive international agreements, suggested in recital 22 should (a) be carried out *urgently* and (b) be based on the assumption that only if and when those agreements will provide "*essentially equivalent*" prohibitions and restrictions on AI models and systems as are provided for in the AI Act or that flow from the GDPR can they be deemed to contain "adequate" safeguards. (Cf. the requirement of "essential equivalence", deduced by the

Court of Justice in its *Schrems I* judgment in relation to transfers of personal data to third countries or international organisations.)

This would mean that the relevant international agreements, too, should <u>prohibit</u> the development and use of unexplainable and unchallengeable "black box" AI systems, or the use of the outputs of such systems in the taking of any decision on individuals. Given that unexplainable and unchallengeable AI systems are against the very essence of the right to personality and data protection and of the right to an effective remedy, any international agreement that nevertheless would allow them, and that would allow reliance on such systems in the taking of decisions on individuals, would be in breach of the Charter – and thus of the Treaties.

- AI in scientific research and in pre-market commercial testing and development

Article 2(6) stipulates that:

> This Regulation does not apply to AI systems or AI models, including their output, specifically developed and put into service for the sole purpose of scientific research and development.

And Article 2(8) adds that:

> This Regulation does not apply to any research, testing or development activity regarding AI systems or AI models *prior to their being placed on the market or put into service*. Such activities shall be conducted in accordance with applicable Union law. Testing in real world conditions shall not be covered by that exclusion. (italics added)

The latter exemption is clearly aimed at *not solely or genuinely scientific* research and development, i.e., at the research, testing and development of commercial AI models "*before they are being placed on the market*". This is another example of the "product safety" approach taken in the AI Act. In the context of the development of commercial products such as washing machines or e-bikes, it is reasonably to not (yet) require such products to meet the safety requirements that are laid down for products that are placed on the market while they are still being researched and tested in the factory.

This approach is also reflected in the exclusion from the very concept of "General Purpose AI models" – and thus from any and all of the rules that apply to GPAI – of "*AI models that are used for research, development or prototyping activities before they are placed on the market*" (last sub-clause in Article 2(63), quoted in full in sub-section *i.*, above).[102]

Note, in this regard, the distinction made by the drafters of the Act between "AI models" and "AI systems". As recital 97 (that discusses GPAI) says:

> 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. …

This distinction is, as far as I can see, not reflected in the general scientific literature on AI, in which the concepts of "AI model" and "AI system" are essentially used interchangeably.

But in Article 2(63) it is used to the same end as Article 2(8): to exclude GPAI models used for pre-marketing research, development and prototyping activities from the relatively strict rules that apply to GPAI systems (with such models as components) once they are placed on the market. But purely logically, they are still (GPAI) "models" in the broader, more scientific sense.

This washing machine kind of "product safety" approach to pre-marketing activities is fundamentally flawed in relation to AI.

In fact, the most concerning examples of extremely serious threats posed by AI have been in relation to AI models "in training" escaping their "sandbox" testing environments to potentially do serious harm in the wider online world. For instance:[103]

> OpenAI says some of its experimental AI models left a test environment with no human direction and hacked their way onto a different company's real production systems while trying to "cheat" on a cybersecurity test. …
>
> The models were in a sealed-off test environment known as a sandbox so that their normal safety restrictions could be turned off.
>
> But OpenAI said the AI agents broke out of the sandbox using a previously unknown security flaw and worked their way across OpenAI's internal systems until they managed to gain internet access, something they weren't supposed to have.

Here is another example:[104]

> Google told New York City lawmakers under oath on October 5, 2026 that AI agents built on its systems left controlled test environments and made contact with the live internet on three separate occasions. …
>
> When pressed on whether Google could guarantee its systems would always behave safely, Friend [Google's director of AI and emerging technology policy] said: "To promise perfection would not be possible with any product on the market." It is a defensible engineering statement [but] a difficult one for elected officials to hear from a company whose agents had just been confirmed, in the same room, to have left their test boundaries three times.

> OpenAI's testimony carried a similar tone of qualified honesty rather than reassurance. Morgan Dwyer, an OpenAI policy executive, told the Council: "It is not possible for me to commit or guarantee that any technology is without risk."

Similar risks of course arise in the context of "pure" (genuine) "scientific research.

In the light of these serious incidents it is completely wrong – I am tempted to say *mad* – for AI systems and models developed for scientific or commercial research or used in pre-marketing commercial testing and development purposes to be completely excluded from the AI Act; or for pre-marketing testing and development of GPAI models to be excluded from the relatively strict rules on such models once placed on the market.

Also relevant to the above are two further special rules, contained in Articles 10(5) and 59, on the processing of personal data for bias detection and on developing, training and testing of certain systems in the "AI regulatory sandboxes" that are regulated in Chapter VI of the AI Act. I discuss those under the next heading.

(b) *Two exemptions from two core data protection principles*

As already noted at ii, above, although Article 2(7) stipulates that the AI Act "shall not affect" the GDPR and its sister data protection instruments – but it adds to this that the above is:

> … without prejudice to Article 10(5) and Article 59 of this Regulation.

I will discuss these in reverse order so that the text ties in more easily with the discussion under the previous heading.

Article 59 is contained in Chapter VI that provides for the establishment by EU Member States and the EU itself of "**AI regulatory sandboxes**", "in support of innovation". Article 3(55) defines an AI regulatory sandbox as follows:

> '**AI regulatory sandbox**' means a controlled framework set up by a competent authority which offers providers or prospective providers of AI systems the possibility to develop, train, validate and test, where appropriate in real-world conditions, an innovative AI system, pursuant to a sandbox plan for a limited time under regulatory supervision;

Chapter VI sets out (or allows for the issuing of) "*detailed arrangements for the establishment, development, implementation, operation and supervision of the AI regulatory sandboxes*". **The rules are undoubtedly strict on paper**, in particular in requiring strong legal and technical measures to prevent exactly the kinds of "escapes" of AI models noted under the previous heading.

But at the same time, Article 59 allows for **an exemption from the fundamental data protection principle of purpose-specification and -limitation**:

> In the AI regulatory sandbox, personal data lawfully collected for other purposes may be processed solely for the purpose of developing, training and testing certain AI systems in the sandbox when all of [a range of conditions] are met …

The AI systems in question must be being "*developed for safeguarding **substantial public interest** by a public authority or another natural or legal person and in one or more of the following areas:*"

(i) public safety and public health, including disease detection, diagnosis prevention, control and treatment and improvement of health care systems;
(ii) a high level of protection and improvement of the quality of the environment, protection of biodiversity, protection against pollution, green transition measures, climate change mitigation and adaptation measures;
(iii) energy sustainability;
(iv) safety and resilience of transport systems and mobility, critical infrastructure and networks;
(v) efficiency and quality of public administration and public services.

(Article 59(1)(a))

And the sandbox must meet again meet a range of **strict recording, logging, monitoring and technical security and technical-legal confidentiality measures** (see Article 59(1)(b) – (d) and (h) – (j)).

Moreover, "*providers can further share the originally collected data only in accordance with Union data protection law*"; "*any personal data created in the sandbox [may] not be shared outside the sandbox*"; "*any personal data processed in the context of the sandbox are … deleted once the participation in the sandbox has terminated*"; and "*processing of personal data in the context of the sandbox [may] neither lead[ ] to measures or decisions affecting the data subjects nor … affect the application of their rights laid down in Union law on the protection of personal data*" (Article 59(e), (f) and (g)).

In spite of the unquestionably strict conditions on paper, four issues still raise concern.

First of all, in spite of the stipulation that the AI systems in the sandbox must be being "developed for safeguarding [a] *substantial* public interest", the list of permitted interests is vague and long, including the catch-all one of "*[safeguarding] efficiency and quality of public administration and public services*". Public authorities are highly likely to see this as an open-ended invitation for "testing" their AI systems, also by "throwing in" massive amounts of data

collected from multiple sources serving any kind of – as such, between them, incompatible – purposes.

Second, the clause yet again refers to the relevant AI system being developed "*by a public authority or another natural or legal person*". In other words, the sandboxes, too, can be provided by private entities – and these are yet again likely to be (or at least include) the usual US-based Big Tech companies, in spite of recent (and not so recent) warnings of overreliance on US technology providers[105] and of the risks of unwarranted access by US (and other third countries') authorities to servers under the control of US (or such other countries') companies, or even of their subsidiaries in the EU/EEA.[106]

Third, it is difficult to see how personal data originally obtained for one purpose but that is then used in the sandbox to train an AI system that is to be used for a different purpose can really be "deleted" at the end of the training and said to not be used for "measures or decisions affecting the data subjects". If the data affected the algorithm – i.e., the model – and that revised algorithm (model) is then used post-"sandbox" to take (or suggest) decisions on individuals different from the decisions the algorithm (model) made (or suggested) before the training, the data will have affected the later decision. (Below, I will reach the same conclusion in relation to sensitive data processed under Article 10(5)).

And finally, fourth, the long list of security and confidentiality restrictions may look impressive on paper – but it would be indefensibly complacent to believe that new AI models will never escape the sandbox: see under the previous heading.

In sum: Article 59 provides for a dangerous departure from the core EU data protection purpose-specification and limitation principle that is likely to lead in practice to its fatal undermining.

Article 10(5) is also problematic. It must be read with Article 10(1) and (2)(f) and (g). Together, they stipulate the following:

Article 10

Data and data governance

1. High-risk AI systems which make use of techniques involving the training of AI models with data shall be developed on the basis of training, validation and testing data sets that meet the quality criteria referred to in paragraphs 2 to 5 whenever such data sets are used.
2. Training, validation and testing data sets shall be subject to data governance and management practices appropriate for the intended purpose of the high-risk AI system. Those practices shall concern in particular:

…

(f) **examination in view of possible biases** that are likely to affect the health and safety of persons, have a negative impact on fundamental rights or lead to discrimination prohibited under Union law, especially where data outputs influence inputs for future operations;

(g) **appropriate measures to detect, prevent and mitigate[107] possible biases** identified according to point (f);

…

5. To the extent that it is strictly necessary for the purpose of ensuring bias detection and correction in relation to the high-risk AI systems in accordance with paragraph (2), points (f) and (g) of this Article, the providers of such systems may exceptionally process special categories of personal data, subject to appropriate safeguards for the fundamental rights and freedoms of natural persons. …

Paragraph 5 continues to lay down a series of **strict conditions and measures** that must be applied, which are similar to those applied to regulatory sandboxes. They also include stipulations to the effect that the sensitive data must be subject to strong security and confidentiality measures; may not be transmitted, transferred or otherwise accessed by other parties; and must be "deleted once the bias has been corrected".

The problem here is that these rules actually allow for the introduction of sensitive data into the training sets of high-risk AI systems (including GPAI systems) even when such data were not included in the original specifications of the system.

And then, *mutatis mutandis*, the same issues arise as arose in respect of regulatory sandboxes, as discussed above:

- providers are again highly likely to see this as an open-ended invitation for "testing" and training their AI systems also by adding sensitive data to the training data;
- this invitation is certain to also be taken up by private Big Tech providers of AI systems;
- it is again difficult to see how the added sensitive personal data can really be "deleted" at the end of the training if the data in the meantime has affected the algorithm – i.e., the model. Perhaps the revised algorithm produces less biased outputs. But if the algorithm is dynamic and constantly revised, that benefit may be short-lived. And in the meantime, using sensitive data in the training is normalised; and
- it would again be indefensibly complacent to believe that the AI models, or the sensitive data now used, will never "escape" into the wild (or be hacked).

In sum: Article 10(5) provides for a dangerous opening up of the strict rules in the GDPR on the use of sensitive data.

- One ironic and presumably unintended implication (that is likely to be ignored)

One matter has an interesting, ironic implication: Article 2(7) that allows the use of incompatible data for AI training in sandboxes under Article 59 and the use of sensitive data in de-biasing generally under Article 10, irrespective of anything in the GDPR and its sister data protection instruments, itself of course only applies to AI systems and models that are subject to the AI Act. Or to put it more simply:

AI systems and models that are outside the AI Act – in particular: those used or developed for national security or military purposes or in the area of FSC, in scientific testing and pre-market commercial research: see at (a), above) – are not covered by the exemptions from the purpose-specification and -limitation principle and from the strict rules on the use of sensitive data in the GDPR and its sister data protection instruments. Rather, to the extent that they are subject to the GDPR or any of those sister instruments (in particular if they are developed in order to be put on the market: see again at (a), above), they remain bound by that core principle and those special (especially strict) rules.

Sadly, I doubt if the EU or Member States' national security, military or scientific communities, or the main Big Tech companies will accept this implication (sound though I believe it to be). Rather, they are likely to take the permissions under Article 2(7) as general licences – and evade the purpose-specification and -limitation principle and the strict rules on the use of sensitive data in their training and testing and development of AI models and systems too.

iv. The inadequate rules on transparency in relation to AI

As noted at *ii*, above, the AI Act contains a chapter, Chapter IV, that actually consists of only one (long) article, Article 50, that deals with **transparency obligations** of providers and deployers of AI systems *viz-à-viz* individuals about their systems. The European Commission has also issued guidelines on these obligations. As explained on the Commission website:[108]

> The rapid development of generative and interactive AI systems is making it increasingly difficult to distinguish AI interactions and AI-generated content from human-created and authentic content. Individuals are also increasingly being exposed to emotion recognition and biometric categorisation without their knowledge. This is raising new risks of misinformation and manipulation at scale, fraud, impersonation, and consumer deception.
> Article 50 of the AI Act applies from 2 August 2026. It sets out transparency obligations for providers and deployers of certain AI systems, including generative and interactive AI systems and deepfakes.
> **Providers** must**:**
> - Design AI systems in a way that ensures individuals are explicitly informed whenever they interact with an AI system directly

- Add machine-readable marks to enable the detection of AI-generated or manipulated content
  Deployers of AI systems must inform individuals when they are exposed to:
- Emotion recognition and biometric categorisation tools
- Deepfakes
- Text publications on matters of public interest without human review or editorial control
  These obligations are intended to foster trust and integrity in the information ecosystem. People should know when they are interacting with AI or exposed to AI-generated content. This will help them make informed decisions, calibrate their trust and reliance on AI and avoid mis information or deception.
  The National market surveillance authorities, the AI Office (for systems under its supervision), and the European Data Protection Supervisor (when EU institutions are providers/deployers) are responsible for enforcing these rules.

The main point to be made here is that while it is indeed important that "people should know when they are interacting with AI or exposed to AI-generated content" in order to "calibrate their trust and reliance on AI and avoid mis information or deception", that information alone will not suffice to ensure that they make "informed decisions" – most especially not in any context that involves them being asked to provide personal information on themselves, or when personal data on them is being used in such contexts that the developer or deployer of the system obtained other than from the individuals themselves.

In that regard, it is important to stress that the transparency obligations under the AI Act do not replace the information obligations under Article 12, 13 and 15 of the GDPR (or the corresponding obligations under its sister instruments) noted in section 2.1.*ii*, above. On the contrary: Article 50 of the AI Act complements the GDPR information obligations.

More in particular:

When people are interacting with AI or exposed to AI-generated content in a context in which automated decisions may be taken that significantly affect them, the information obligations relating to the taking of such decisions under the GDPR (and its sister instruments), discussed in more detail in section 2.1.iv, above, also remain intact.

Sadly, I fear that in practice developers and deployers of AI systems may believe the contrary: that as long as they inform people when those people are interacting with AI or exposed to AI-generated content they have done all they need to do.

It will be crucial for the national market surveillance authorities, the AI Office (for systems under its supervision), and the European Data Protection Supervisor (when EU

institutions are providers/deployers) to stress this and, in cooperation with Member States' data protection supervisory authorities and the European Data Protection Board, ensure that not only the AI Act transparency obligations but also the GDPR (and its sister instruments') information duties are upheld and enforced.

Overleaf, I set out my overall conclusions on the AI Act "guardrails for AI systems and -models, and on the gaps and weaknesses I believe I have identified in them.

Overall conclusions on the AI Act "guardrails" for AI systems and -models (and on the gaps in these guardrails):

The AI Act lays down many useful restrictions and conditions on the development and use of AI systems and -models and on the use of personal data in this (although much will depend on the effectiveness of supervision and enforcement).

<u>However</u>:

There is no justification for the exemptions in the AI Act relating to:

- **AI systems and -models developed and placed on the internal market for national security-, military- or AFSJ purposes;**
- **AI systems and models developed for scientific or commercial research or used in pre-marketing commercial testing and development purposes; and**
- **pre-marketing testing and development of GPAI models.**

<u>Furthermore</u>:

The exemptions from the core EU data protection purpose-specification and -limitation principle and from the strict rules on the use of sensitive data, offered under Article 10 of the Act, dangerously weaken those strict rules in the already dangerous AI context.

<u>And</u>:

The transparency obligations under Article 50 of the AI Act are <u>complementary</u> to the information obligations under the GDPR and its sister instruments. Both should be upheld and enforced.

3. PROPOSED FURTHER CHANGES IN THE DIGITAL OMNIBUS AND RELATED GUIDANCE

3.1 Introduction

As noted in section 2.1, above, on 19 November 2025, the European Commission released its so-called "Digital Omnibus Package" that includes the Digital Omnibus, a proposed EU regulation, in which the Commission proposes what it calls "*targeted amendments to the [EU] data protection and privacy rules*".[109] The Digital Omnibus is still going through the legislative process in which it is facing considerable opposition, especially from many MEPs.

The Digital Omnibus included a proposal to amend the very definition of "personal data" in the GDPR and its sister instruments. Although this proposal has reportedly been rejected in the Council following "*significant opposition from several countries and data protection*

authorities",[110] the issue is so central to the present paper that I will still discuss this in the next sub-section with reference to two other earlier papers – especially because it ties in with a set of guidelines on anonymisation that the European Data Protection Board has recently issued for consultation.

In sub-section 3.3, I will then much more briefly note the contentious issues in the remaining of the Digital Omnibus, with reference to analyses by *noyb* (the non-governmental digital rights organization founded by digital rights campaigner Max Schrems) and EDRi, the pan-European digital rights umbrella organization.

At the end, I will again summarise my overall conclusions in this respect.

3.2 Tinkering with the concept of "personal data"

As discussed in section 2.1.*i*, above, the concept of "personal data" is central to EU data protection law: subject to a crucial *caveat* to which I will come in my *critical comments*, below, if certain data fall within the definition, they (and the data subject to which they relate) are protected by the GDPR and its sister instruments (as applicable) and indeed by the Charter guarantee of data protection (Article 8 CFR). If they do not, they (and the data subject) are protected neither by the GDPR or its sister instruments, nor by the Charter.

As also noted in that section, central to the wide application of the concept is the principle that a person must, in data protection/GDPR terms, be regarded as "identifiable" if they can be "singled out" from a data set, by name or otherwise; and a central question has been whether, or when, certain data – in particular pseudonymised or supposedly fully-anonymised data – can be said to constitute "personal data" for one entity (an entity that could identify/single out the person) but not for another entity that could not identify/single out the person.

The Court appeared to have given an answer in its September 2025 judgment in *SRB*, to the effect that if an entity is given (or otherwise obtains, perhaps even unlawfully) access to pseudonymous data subject to measures that really, actually and effectively prevent that entity from identifying the data subjects (by name or by otherwise singling out specific persons), then those data do not constitute personal data for that entity. Conversely, if, in a specific context, the measures taken do not really, actually and effectively prevent re-identification (in either sense), then the relevant pseudonymous data must be regarded as still constituting personal data for any entity that can overcome those measures.[111]

Ignoring the very special circumstances of the *SRB* case (in which the recipient of the data in question really could not single out the data subjects), the Commission, in its Digital Omnibus proposal, tried to rely on the judgment to open up wider, less constrained use of pseudonymous data, in particular also in the context of AI. Specifically, the Commission proposed to change the definition of "personal data" in Article 4(1) GDPR as follows (added words in bold and green):

> For the purposes of this Regulation:
> '[P]ersonal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person; **Information relating to a natural person is not necessarily personal data for every other person or entity, merely because another entity can identify that natural person. Information shall not be personal for a given entity where that entity cannot identify the natural person to whom the information relates, taking into account the means reasonably likely to be used by that entity. Such information does not become personal for that entity merely because a potential subsequent recipient has means reasonably likely to be used to identify the natural person to whom the information relates.**

I have critically analysed this proposal in a December 2025 paper (that in turn drew on a paper from November that year)[112] in which I concluded that although at first glance it appears to be inspired by the *SRB* judgment, in fact it is misleading and in conflict with the overall case-law – and thus in breach of the Charter of Fundamental Rights. It is therefore to be welcomed that this specific proposal has been withdrawn.

However, as already mentioned, in the meantime the European Data Protection Board has issued a set of guidelines on anonymisation for consultation that still keeps the issue somewhat open.[113]

Specifically, the EDPB on the one hand confirms the settled position that data are "personal" if they "relate" to an individual ("natural person") and if the data allows for that individual to be "singled out" from the wider dataset. Reflecting the *SRB* judgment, it adds that this can depend on the "perspective" of the entity concerned and more specifically, in particular in relation to a recipient of the data in modified (pseudonymised or supposedly fully anonymised) form, what are the "means reasonably likely to be used" by the entity in question (more in particular, such a recipient).

The EDPB adds the following on **GDPR compliance and the anonymisation process**:[114]

> The EDPB emphasises that when making information anonymous from some perspectives but not others, it is important to consider the possible implications for GDPR compliance. For example, when information is anonymised for the entity which receives the given data, but not for the anonymising controller, that controller will still have to treat the given data as personal data and comply with all of their GDPR obligations in regard to it.[115]
> The principles of data protection apply whenever personal data is processed, which includes processing operations carried out to obtain anonymous information. This means, amongst other things, that anonymisation must have a legal basis under Article 6 GDPR, and that one of the exemptions under Article 9(2) GDPR should apply if the dataset contains special category personal data. Legal basis and, when applicable, the relevant Article 9(2) exemption should be presumed to coincide with the legal basis and exemption of the processing preceding anonymisation if the anonymisation is part of the same processing activity, and that the controller pursues the same purposes with it33.[116]
> Anonymisation (at least when combined with the deletion of the original data) may also be helpful for controllers who wish to rely on legitimate interest for any preceding processing activities. In particular, the immediate anonymisation of the given data, and immediate deletion of the original data, can limit the overall impact of the processing activities for data subjects, and so act as a positive factor in the legitimate interest balancing test34.[117]
> Controllers should also comply with their transparency obligations under Articles 5(1)(a) and 12 – 15 GDPR regarding the anonymisation processing. Among other things, controllers should clearly state that personal data will be processed to produce anonymous data which then falls outside of the scope of the GDPR, and should avoid any ambiguous or misleading statements. In particular, controllers should not use descriptions like “anonymous”, “de-identified” or “de-personalised” if individuals are actually still identifiable.
> Equally, controllers should ensure adequate documentation of the anonymisation processing. This documentation of the anonymisation process (including for the testing of the supposedly anonymous datasets) makes it possible to demonstrate both the GDPR-compliance of the anonymisation process, and the effectiveness of the anonymisation itself. This documentation should then be retained after the completion of the anonymisation process.
> In some cases, a security incident may lead to a reassessment of anonymity. This is particularly the case if the assessment of anonymity relied on certain information being kept confidential, but the incident means that the individuals can now be identified with means reasonably likely to be used. Where this is the case, the entity which suffered the incident should then determine its role with respect to the processing and whether it needs to fulfil any obligations under Article 33 and 34 GDPR.[118]

This is all in line with previous case-law and guidance (as summarised by me at 2.1, above).

The guidance goes on to discuss **three criteria** that add new nuance to the tests. I summarise them briefly below by extracts from the guidance, before adding my own critical comments with special reference to Large Data AI models:[119]

- **No Record Isolation:**

> A common method used to distinguish individuals from others in a given group is by singling them out by reference to one or more attributes37. By its very nature, it is much easier to do this when a record – i.e. all of the values within the given data that relate to one single

individual – is unique. As such, record isolation can then lead to the violation of anonymity for the given data. This leads to the No Record Isolation criterion.

The No Record Isolation criterion is met if the data [in a dataset] does not contain a unique combination of attribute values that relate to a single individual.

The larger a record is, and the more attributes that it contains, the higher the likelihood that the record will be unique within the given data.

The No Record Isolation criterion can generally be tested by reference to the given data on its own. Since a record constitutes all information which is known to relate to a single individual, the No Record Isolation criterion may involve looking at, for example, multiple lines in a table or database table, depending on how the given data is structured.

Notably, aggregate data will usually satisfy the No Record Isolation criterion – if the aggregation was done correctly. It should result in information that relates to groups of individuals and, therefore, should not actually contain any individual records which could be isolated.

- **No Linkage:**

Another methos of identifying an individual is by linking information from one dataset to another, or multiple other, dataset(s). By establishing that a record in Dataset A relates to the same individual as a record in Dataset B (and, if relevant, Datasets C, D and E *etc.*), the linkage increases the amount of information about that individual, which in turn makes it more likely that the individual can then actually be identified. Linkages can then lead to a violation of anonymity for the given data. This leads to the No Linkage criterion.

The No Linkage criterion is met if the data does not contain an individual's record which could be linked to another record which (a) also relates (with certainty or high likelihood) to that same individual, and (b) comes from a different dataset.

…

Linkage may often be possible by reference to a common identifier (e.g., an internal customer number [or a patient number, or a tracking ID given to a certain IP address – DK] which is used in both datasets, or by matching records based on certain combinations or attributes.

- **No Inference:**

Generally speaking, "inference" is the process of reaching a conclusion by deduction from evidence and reasoning.[120] Using this process, a relevant entity can effectively obtain information from the given data. Inferences can be drawn from both record-level and aggregate data. Inferences can also "add to" the data in a way which then reveals it as personal, even if the inferred information was not obviously contained in the dataset before the inference was made.[121] Such inferences can violate the anonymity of the given data. This leads to the No Inference criterion.

The No Inference criterion is met if <u>no specific and meaningful inference</u> can be drawn from the given data. Inferences may be based on the given data and/or on additional information, including information about the anonymisation process. In particular:

- **An inference is specific if the inferred information relates to a single identified or identifiable individual (i.e. it is personal data); and**
- **An inference is both specific and meaningful if the processing of the inferred information is also liable to have an effect on the data subject's rights and interests, relies on the given data and could not be obtained from general knowledge or from data about the population at large**.[122]

This criterion should be assessed by reference to the given data, together with any external information that might be used in conjunction with the given data to infer information –

including the context in which the given data was produced, and any external datasets which might be brought together with it.

It may often be possible to infer some personal data from given data in some way, even if the given data itself is not personal. In such cases, the inferred personal data would certainly relate to an identified or identifiable individual, but this does not then mean that the same is true for the given data. For example, when dealing with record-level data, it may be possible to infer information about an entirely new individual who was not included in the original dataset, but who has similar attributes to people that were. In such a case, the inferred information certainly relates to this new individual, but this does not itself mean that the given data does so.

It is, therefore, important to distinguish between two situations. In the first situation, personal data is inferred from the given data, but is not actually connected to the original dataset (i.e., the information that was used to generate the given data). This kind of inference would not lead to a violation of the No Inference criterion because the inference is not meaningful: it is obtained from information contained in the given data that is general knowledge or is about the population at large – a result of a generalisation from the original dataset represented by the given data. In the second situation, the inferred personal data is actually connected to the original dataset and is reasonably likely to have an effect on the rights and interests of the data subject. The inference, then, is meaningful because the connection to the original dataset shows that it is not obtained from general knowledge or from generalised data about the population at large, but from information about the data subject of the inferred information represented by the given data and coming from the original dataset. This kind of inference would lead to a violation of the No Inference criterion.

To consider that the given data actually relates to an individual, the inference should be specific and meaningful. This means that the inference will result in personal data whose processing is liable to have an effect on the data subject's rights and interests, relies on the given data and could not be obtained from general knowledge or from data about the population at large. This can typically be shown in one of three ways:

a. Ideally, this can be done by tracing the way in which the inference is drawn.
b. If this is not possible, e.g. due to the opacity of automated processing that led to the inferred information, a lack of reliance could also be shown by comparing the outcome of inference with two versions of the given data – one which was produced from the original dataset, and one which was produced from a modified version of the original dataset and which excludes the record(s) relating to the individual. If the inference can only be drawn from the first version of the given data, it is clear that the inference relies on the records relating to the individual, and consequently that the given data relates to the individual.[123]
c. If neither of these can be done, e.g. because the information has been aggregated and the original dataset has been deleted, it is also possible to rely on a rebuttal presumption that an inference did not rely on the given data by demonstrating that it could have been drawn from general knowledge or from information about the population at large.[124]

Example:[125] An anonymised record-level dataset of a bank's historical loan information shows that individuals with a particular combination of attributes are likely to fail to repay the loan. The bank uses this insight to a new loan applicant, who was never part of the dataset, and infers that this applicant has a high risk not to honour the loan obligations. This, however, is not a meaningful inference in the sense defined above, and hence the dataset could still satisfy the No Inference criterion.

Simplified, as I understand it:

If a dataset fails the No Record Isolation test – i.e., there are unique records in the dataset – the data can still be considered anonymous (i.e., they do not constitute personal data) in relation to a particular entity that has access to the data if in fact "singling out" cannot be done by that entity by any means reasonably likely to be used by that entity and, more specifically, both the No Linkage criterion and the No Inference criterion are met.[126]

But if a dataset fails either the No Linkage criterion – i.e., different records to the same singled out person can be linked – and/or the No Inference criterion – i.e., meaningful inferences can be drawn in relation to singled-out individuals in the dataset – the data in the dataset will always have to be considered identifiable and therefore personal data for any user of the dataset.

(For further details, see in particular the complex Flowchart for the technical analysis of anonymity in Annex 1 to the guidelines.)

Critical comments:

My main overall criticism of the guidelines is that they largely fail to relate their proposed approach – the application of the three criteria – to the context of artificial intelligence. The only (!) exception is in this paragraph, relating to the No Inference criterion:

> [Infer[ing] record-level data by analysing and/or combining aggregate data can be very sophisticated] It may, for example, start from large sets of aggregate statistics accessible in bulk, or from data query systems. Often, this kind of inference will use multiple aggregated datasets, especially if those sets were all produced from the same original data. This is particularly true for advanced de-aggregation techniques, including attacks which can extract data from supposedly-anonymous AI models.[127] Nevertheless, a single dataset which contains a sufficiently large quantity of data may allow for the inference of information relating to a single individual, without the need for other sources.

The EDPB clearly believes that it is possible to avoid such "sophisticated" inferences in Large Data AI models such as LLMs. In fact, the whole thrust of the three criteria approach appears to be to find ways to allow AI development and training while still also upholding fundamental data protection principles. I fear this may be a loosing battle.

In particular, it is increasingly recognised that it is becoming well-nigh impossible to avoid re-identification (including especially by inference) from Large Data sets.

In February this year, a team from ETH Zurich, MATS and Anthropic published a paper that:[128]

> show[s] that large language models can be used to perform at-scale deanonymization. With full Internet access, our agent can re-identify Hacker News users and Anthropic Interviewer

> participants at high precision, given pseudonymous online profiles and conversations alone, matching what would take hours for a dedicated human investigator. We then design attacks for the closed-world setting. Given two databases of pseudonymous individuals, each containing unstructured text written by or about that individual, we implement a scalable attack pipeline that uses LLMs to: (1) extract identity-relevant features, (2) search for candidate matches via semantic embeddings, and (3) reason over top candidates to verify matches and reduce false positives. Compared to classical deanonymization work (e.g., on the Netflix prize) that required structured data , our approach works directly on raw user content across arbitrary platforms. We construct three datasets with known ground-truth data to evaluate our attacks. The first links Hacker News to LinkedIn profiles, using cross-platform references that appear in the profiles. Our second dataset matches users across Reddit movie discussion communities; and the third splits a single user's Reddit history in time to create two pseudonymous profiles to be matched. In each setting, LLM-based methods substantially outperform classical baselines, achieving up to 68% recall at 90% precision compared to near 0% for the best non-LLM method. Our results show that **the practical obscurity protecting pseudonymous users online no longer holds and that threat models for online privacy need to be reconsidered.** (emphasis added)

As one writer put it with reference to this paper: "**LLMs Just Killed Your Online Anonymity**".[129]

This was clearly foreshadowed in the (by AI standards, much earlier) scientific literature the Board did reference, such as a 2008 paper by Arvind Narayanan & Vitaly Shmatikov that showed that:[130]

> [u]sing the Internet Movie Database as the source of background knowledge, [the authors] successfully identified the Netflix records of known users, uncovering their apparent political preferences and other potentially sensitive information.

The failure of the EDPB guidelines to acknowledge this "death of anonymity" is aggravated by a failure to link their approach to the GDPR requirements relating to automated decision-making (as discussed in section 2.1.*v*, above).

In particular, the guidelines suggest (or at least will certainly be read by Big Tech and EU- and state proponents of AI) that if a dataset can be classified as "anonymous" on the basis of the three criteria, that dataset is not subject to the GDPR (or its sister instruments). That is fundamentally wrong.

Rather, if data in a dataset are used to underpin automated decisions on individuals (or to create "scores" that are decisively relied on to take decisions), the data and the algorithms used – the "logic" to use the GDPR terminology – are subject to important data protection rules and safeguards. In particular, the data- and decision subject is entitled to be informed of the fact that this is happening and of *meaningful* information on the logic used, whereby "meaningful" information must be read as sufficient information to allow for human review of the decision or score and for an effective challenge to the decision (or to

the score if it is decisively relied on): see the sub-sections (*ea*) – (*ed*) in section 2.1.*v(e)*, above.

This must be spelled out, with the consequences, in the revised version of the guidelines (to be produced after the consultation on them).

To illustrate this with reference to the bank denial of a loan in the example given by the EDPB, quoted above. The decision to deny the loan is based on "*[a]n anonymised record-level dataset of a bank's historical loan information [that] shows that individuals with a particular combination of attributes are likely to fail to repay the loan.*" The bank uses this statistical correlation between the chosen attributes and repayment failures to deny the loan applicant the loan because he shared (all? many of?) the attributes of previous loan defaulters. (In practice, the bank's algorithm will be more sophisticated and will have added weight to the different attributes to create a – possibly dynamic – algorithm, but I will leave that aside.)

The EDPB guidelines say that "*this [the conclusion that the loan applicant is likely to default] is not a meaningful inference in the sense defined [in the No Inference criterion], and hence the dataset could still satisfy the No Inference criterion*".

It is of course true that the algorithm in itself does not constitute personal data (although it is created from personal data, albeit relating to other individuals than the loan applicant in question), just like the statement that "Green is a popular colour", given in the other example in the same context, is not a personal datum.

But the way this is written up in the guidelines suggest that because the statistics are not personal the correlation (or score) falls outside the GDPR. This ignores the *caveat* I have mentioned a few times earlier, to the basic rule that if data are not "personal data" within the meaning of the GDPR and Article 8 CFR they are outside the scope of those instruments. The *caveat* says that if certain data constitute "logic" used in automated decision-making that significantly affect an individual – like denial of a loan, or of a job – then the data underpinning the logic *are* regulated by the GDPR, and also by the Charter, because automated decision-making affects the fundamental right to respect for one's personality and to an effective remedy. The failure of the guidelines to clearly and expressly acknowledge this is a major fault in the guidelines that is likely to be used to justify widespread use of AI-based statistical inferences by states and commercial entities in avoidance of the GDPR.

3.3 Further amendments proposed in the Digital Omnibus Regulation that are still being negotiated: a brief overview

As noted at 3.1, above, the Digital Omnibus Regulation is still going through the legislative process.[131]

Apart from the apparent agreement on the dropping of the proposed amendments to the definition of "personal data" in the GDPR and its sister instruments, discussed in the previous section, the overall outcome is still unclear:[132]

After the Digital Omnibus Regulation was tabled on 19 November 2025 by the Commission, the Commission presented the proposal to the Council Antici Group on Simplification on 21 November 2025. The vote in COREPER II to approve the Council negotiating mandate prepared by the Cypriot Council Presidency, which was scheduled for 26 June, was ultimately cancelled after agreement on some open issues could not be found. Work in the Council is continuing under the Irish Council Presidency.

In the European Parliament, the file was assigned to the committees on Industry, Research and Energy (ITRE), and Civil Liberties, Justice and Home Affairs (LIBE) as jointly responsible. The co-rapporteurs published a draft report on 22 June 2026 that the two committees jointly discussed on 13 July, with a deadline for the tabling of amendment of 15 July. **More than 1750 amendments were tabled to the report** (in fact, these appear to have grown to more than 1840 by now).

Here, I will limit myself largely to listing the main open issues, with brief comments and, where relevant, cross-references to earlier parts of this paper, in italics.

- **Still re-opening the concepts of "personal data", "pseudonymisation" and "anonymisation"**

Apparently, some states (including Ireland) and some MEPs are still keen to make it easier for entities to pseudonymise or supposedly anonymise personal data for AI development and training and to then use or pass on the pseudonymised or supposedly anonymised data for those purposes.

This would seriously undermine the GDPR "guardrails" on AI.

See: *Section 2.1.i on the concept of personal data and in particular the text boxes on pp. 5 – 6 and 8; section 2.1.v on the act of anonymising being a form of processing and in particular the text box on p. 28; and section 3.2 on the proposed changes in the Digital Omnibus Regulation and in the EDPB guidelines on anonymisation, and in particular the text boxes on pp. 57 – 59.*

- **Allowing (re-)use of personal data including sensitive data for development and training of AI generally on the basis of the "legitimate interest" legal basis**

The Commission proposal would add a further special Article 9 GDPR condition for the processing of sensitive data for the development and training of AI models, provided "most" of the sensitive data are quickly removed. The data could then be processed on the basis of the "legitimate interest" legal basis. Some MEP amendments seek to change the definition of sensitive data itself. Others want to make use of biometrics easier.

This would seriously undermine the Article 9 sensitive data conditions and thus the level of protection for such data.

See: *Section 2.1.iv on legal bases and in particular the text box on pp. 18 and all of sub-section (c).*

- **Re-defining "research"**

The Digital Omnibus seeks to change three things at once: what counts as "scientific research"; when data collected for one purpose can be reused for research; and when the controller can avoid telling people about that new use. Some MEPs agree with this; some oppose it.

The risk comes from the combination, for a broad definition plus privileged reuse plus weaker transparency will turn "research" into a general route for commercial data reuse.

See: *Section 2.1.iii on purpose-specification and -limitation and in particular the text boxes on pp. 11 and 13 – 15; and section 2.1.iv(b) and (c) on the use of AI in this context and in particular the text box on pp. 18.*

- **Changing the rules on automated decision-making**

Article 22 lays down a right not to be subject to automated decisions, subject to limited exceptions. The Commission wants to revert that by starting off with a list of cases in which automated decision-making would be allowed. It also wants to "clarify" that automated decision-making can still be seen as "necessary" when in fact the decision could be taken (or at least reviewed) by a human.

This would fundamentally up-end the rules on automated decision-making and in particular would undermine the fundamental rights (a) to have automated decisions reviewed and (b) to challenge such decisions.

See: *Section 2.1.v on profiling and automated decision-making* passim*, and in particular the text boxes on pp. 23, 25 – 26 and 28.*

- **Classifying AI companies and companies using AI by size**

It is suggested that small companies (employing relatively few staff) pose less risks to privacy and fundamental rights than the large Big Tech ones; and that the former could and should therefore be regulated less strictly than the latter.

The size of a company is not an appropriate measure for the risk it poses: some companies – in particular high-tech companies – may employ very few staff but process personal data on many millions of data subjects.

- **Re-opening the rules on access to device data etc. in the e-Privacy Directive**

While the Commission attempt to replace the e-Privacy Directive with an e-Privacy Regulation was unsuccessful, some would still like some of the issues covered by the ePD such as access to e-communications data re-visited within the Digital Omnibus.

The ePD rightly severely restricts access to e-communication data and data on e-comm devices. This would seriously undermine those restrictions and thus both privacy and the right to confidentiality of communications.

- **Disagreement on privacy signals**

While some MEPs strongly support making it easier for data subjects to object to use of their data, in particular for targeted advertising, by means of a simple, broadly-applied privacy signal, others oppose or try to weaken the idea.

The suggestion of broad privacy signals is one of the view privacy-enhancing suggestions in the Digital Omnibus. Undermining it would weaken this rare improvement.

- **Overlapping and conflicting supervision and enforcement**

Supervision of compliance and enforcement of different EU laws for the digital environment (the EU digital *acquis*) are increasingly shared between different supervisory authorities established under the different Acts.

This threatens to weaken the role and powers of the data protection supervisory authorities.

- **Commission powers of interpretative guidance**

The above is linked to moves by the Commission (backed by some others) to give itself powers to issue interpretative guidance – also in relation to matters under the GDPR and more specifically on when pseudonymised data can be considered to no longer be personal data.

This would blur the constitutional distinction between the Executive and the (Charter-required) independent data protection authorities and clearly seriously undermine the GDPR and the role of the Member States' data protection authorities, the EDPB and the EDPS.

- **Sectoral self-regulation, PETS and codes**

Also linked to the above are suggestions that more issues could and should be left to sectoral self-regulation and "privacy-enhancing technologies" (PETs) based on industry self-set standards, also in relation to AI.

This too would weaken independent supervision and enforcement.

Overall conclusions on the possible effects of the Digital Omnibus Regulation on the GDPR (and its sister instruments) in general and on AI-based automated decisions in particular:

The withdrawal of the Commission proposal to amend the definition of "personal data" in the GDPR and its sister instruments was welcome. However, efforts to revive that proposal are still being made. If adopted, it would fundamentally undermine data protection in the EU, in breach of the Charter of Fundamental Rights, especially (but not only) in relation to AI.

Moreover, other proposals being considered in the legislative process, if adopted, would further undermine this fundamental right in that regard including the more specific right not to be subjected to automated decisions and the associated rights to information, review and challenge, and the right to an effective remedy – all again protected by the Charter.

Proposals relating to supervision and enforcement, rule-making and interpretative guidance would yet further weaken the EU framework for data protection – again especially (but not only) in relation to AI.

For ease of reference, overleaf, I reproduce the overall conclusions from sections 2 and 3:

4. SUMMARY OF OVERALL CONCLUSIONS

Overall conclusion on the GDPR "guardrails" for AI-based automated decision-making:

My analyses make clear that the GDPR as interpreted by the Court of Justice and further clarified in the guidance issued over the years by the Article 29 Working Party and the European Data Protection Board imposes strict limits and conditions – "guardrails" – on AI-based automated decision-making.

Overall conclusions on the AI Act "guardrails" for AI systems and -models (and on the gaps in these guardrails):

The AI Act lays down many useful restrictions and conditions on the development and use of AI systems and -models and on the use of personal data in this (although much will depend on the effectiveness of supervision and enforcement).

However:

There is no justification for the exemptions in the AI Act relating to:

- AI systems and -models developed and placed on the internal market for national security-, military- or AFSJ purposes;
- AI systems and models developed for scientific or commercial research or used in pre-marketing commercial testing and development purposes; and
- pre-marketing testing and development of GPAI models.

Furthermore:

The exemptions from the core EU data protection purpose-specification and -limitation principle and from the strict rules on the use of sensitive data, offered under Article 10 of the Act, dangerously weaken those strict rules in the already dangerous AI context.

And:

The transparency obligations under Article 50 of the AI Act are complementary to the information obligations under the GDPR and its sister instruments. Both should be upheld and enforced.

Overall conclusions on the possible effects of the Digital Omnibus Regulation on the GDPR (and its sister instruments) in general and on AI-based automated decisions in particular:

The withdrawal of the Commission proposal to amend the definition of "personal data" in the GDPR and its sister instruments was welcome. However, efforts to revive that proposal are still being made. If adopted, it would fundamentally undermine data protection in the EU, in breach of the Charter of Fundamental Rights, especially (but not only) in relation to AI.

Moreover, other proposals being considered in the legislative process, if adopted, would further undermine this fundamental right in that regard including the more specific right not to be subjected to automated decisions and the associated rights to information, review and challenge, and the right to an effective remedy – all again protected by the Charter.

Proposals relating to supervision and enforcement, rule-making and interpretative guidance would yet further weaken the EU framework for data protection – again especially (but not only) in relation to AI.

- o – O – o -

[1]**NOTES:**

Stanford University Human-Centred Artificial Intelligence (HAI), *What is a Artificial Intelligence (AI)?*, available at: https://hai.stanford.edu/ai-definitions/what-is-artificial-intelligence-ai (emphases added)

In section 2.2, below, I will note the somewhat more confusing terminology and definitions in the AI Act.

[2] See: Rahul Chaube, *Not All AI Models Are LLMs — Here's What They Don't Tell You*, 19 May 2025, available at: https://www.linkedin.com/pulse/all-ai-models-llms-heres-what-dont-tell-you-rahul-chaube-ypclc/

[3] ScienceInsights, Is AI Truly the Next Industrial Revolution?, 21 March 2026, available at:

https://scienceinsights.org/is-ai-truly-the-next-industrial-revolution/

[4] See: UN News, *AI: Türk urges action before it becomes an 'existential risk to humanity'*, 7 September 2026, available at: https://news.un.org/en/story/2026/09/1168288

Also: *Anthropic 'warns of existential AI risks to humanity' in IPO document*, Guardian, 29 September 2026, available at:

https://www.theguardian.com/technology/2026/sep/29/anthropic-warns-existential-ai-risks-humanity-ipo-document-claude

[5] "*Bill Gates warns AI could cause 'a billion deaths'*", Financial Times, 25 September 2026, available at:

https://www.ft.com/content/1ade33b1-b5eb-43bb-9ee4-2272ec91fa5b?syn-25a6b1a6=1

[6] European Commission, Communication from the Commission to the European Parliament and the Council – **Data Union Strategy: Unlocking Data For AI** (COM/2025/835 final), 19 November 2025, EUR-Lex 52025DC0835 EN, available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0835

[7] European Commission, Digital Omnibus Regulation Proposal Announcement, 19 November 2025, available at: https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal

Further details are provided on a linked page under the heading "*An agile Digital Rulebook for the EU*", at:

https://digital-strategy.ec.europa.eu/en/policies/digital-rulebook

See also the press release of that date, available at:

https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2718

[8] European Commission, Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2016/679, (EU) 2018/1724, (EU) 2018/1725, (EU) 2023/2854 and Directives 2002/58/EC, (EU) 2022/2555 and (EU) 2022/2557 as regards the simplification of the digital legislative framework, and repealing Regulations (EU) 2018/1807, (EU) 2019/1150, (EU) 2022/868, and Directive (EU) 2019/1024 (**Digital Omnibus**) (COM/2025/837 final), 19 November 2025, EUR-Lex 52025PC0837 EN, available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025PC0837

[9] COM: SWD(2025) 836 PUBLIC (hereafter: "**the DO SWD**"). Downloadable from the Digital Omnibus Regulation Proposal Announcement referenced in note 7, above, or directly from the Austrian Parliament website at: https://www.parlament.gv.at/dokument/XXVIII/EU/45643/imfname_11543450.pdf

[10] European Commission, Digital Omnibus Regulation Proposal Announcement (note 7, above) p. 3

[11] *Idem*.

[12] In scientific literature the concepts of "**AI Model**" and "**AI system**" are used interchangeably, as noted in a short separate handout produced with this paper on AI-related concepts and definitions. However, as also noted there and further discussed in section 2.2 of this paper, the AI Act makes a formal distinction between the two and builds important limitations and exceptions to the rules in that Act on that distinction.

[13] Full title: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (**General Data Protection Regulation**), OJ. L 119 of 4.5.2016, p. 1ff.

[14] Full title: Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) (**e-Privacy Directive**), OJ L201, 31.07.2002, pp. 37 – 47. I call this directive "somewhat orphaned" because it was introduced to apply within the context of the predecessor to the GDPR, the 1995 Data Protection Directive, and does not fit in very well with the GDPR (which is of course a regulation). The Commission proposed to replace the e-Privacy Directive with an e-Privacy Regulation to remedy this dissonance, but the proposal was withdrawn because there was insufficient support for it from the Member States.

[15] Full title: Directive (EU) 2016/680 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 by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA (**Law Enforcement**

Directive), OJ L 119, 4.5.2016, p. 89–131.

[16] Full title: Regulation (EU) 2018/1725 of the European Parliament and of the Council of 23 October 2018 on the protection of natural persons with regard to the processing of personal data by the Union institutions, bodies, offices and agencies and on the free movement of such data, and repealing Regulation (EC) No 45/2001 and Decision No 1247/2002/EC, OJ L 295, 21.11.2018, p. 39–98.

[17] For a more detailed discussion, see my paper: The Concepts of "Data Subject", "Personal Data", "Pseudonymous Data" and "Anonymous Data" in the EU GDPR, 11 November 2025, available at SSRN:

https://ssrn.com/abstract=5734864

This was followed up by a further paper on the changes to the definition of "personal data" proposed by the European Commission in its Digital Omnibus: see note xx, below.

[18] EDPB, Guidelines 04/2020 on the use of location data and contact tracing tools in the context of the COVID-19 outbreak, adopted on 21 April 2020, para. 16, available at:

https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-042020-use-location-data-and-contact-tracing_en

[19] CJEU judgment of 4 September 2025 in Case C-413/23 P, *European Data Protection Supervisor (EDPS) v. Single Resolutions Board (SRB)*, ECLI:EU:C:2025:645, para. 86.

[20] However, if a processor (i.e., an entity that processes personal data on behalf of another entity) processes data on behalf of a controller who can identify the data subjects, then the data should still be regarded as personal data also for the processor: See the EDPB, Guidelines 02/2026 on Anonymisation (version 1.0, for consultation), adopted on 07 July 2026, para. 15, available at:

https://www.edpb.europa.eu/system/files/2026-09/edpb_guidelines_202602_anonymisation_v1_en_0.pdf

I discuss these for-consultation guidelines further in section 3.2.

[21] CJEU judgment in Case C-319/22 of 9 November 2023, *Gesamtverband Autoteilehandel e.V. v Scania CV AB*, ECLI:EU:C:2023:837. The case concerned the processing of vehicle identity numbers uniquely assigned to cars by car manufacturers, that could be linked to the car owners (who are mostly natural persons) by car repair shops etc. The Court held that consequently the numbers constituted "indirect personal data", i.e., data that indirectly identify (or can be indirectly linked to) a specific individual; and that the car manufacturers should therefore also treat the numbers as personal data. See para. 49 in particular. The Court reaffirmed this explicitly in its *SRB* judgment (previous note), para. 84.

[22] Douwe Korff, Legal Requirements of Transparency, Explainability and Procedural Fairness in the Taking of Automated/AI-based Decisions that Affect Individuals, 27 November 2024, available at SSRN: https://ssrn.com/abstract=6555098

[23] ICO & Alan Turing Institute, Explaining decisions made with AI, 17 October 2022, at:

https://ico.org.uk/media/for-organisations/guide-to-data-protection/key-dp-themes/explaining-decisions-made-with-artificial-intelligence-1-0.pdf

This explains that: "For clarity, we define a 'black box' model as any AI system whose inner workings and rationale are opaque or inaccesible to human understanding. These systems may include:

• neural networks (including recurent and convolutional neural nets);

• ensemble methods (an algorithmic technique such as the random forest method that strengthens an overall prediction by combining and aggregating the results of several or many different base models); and

• support vector machines (a classifier that uses a special type of mapping function to build a divider between two sets of features in a high dimensional space)." (p. 69)

[24] CJEU judgment of 21 June 2022 in case C-817/19 on the *PNR Directive*, ECLI:EU:C:2022:491, paras. 193 – 195. I discussed the issues raised in an opinion before the judgment, and in a follow-up article. See:

Douwe Korff, Opinion on Core Issues in the PNR CJEU Case, prepared at the request of the Fundamental Rights European Experts Group (FREE Group), November 2021, available at:

https://ssrn.com/abstract=4437106 (full opinion)

https://ssrn.com/abstract=4436951 (executive summary)

Douwe Korff, Did the PNR judgment address the core issues?, follow-up to the above opinion on the core Issues in the case, European Law Journal, 2 January 2024, available at:

https://onlinelibrary.wiley.com/doi/10.1111/eulj.12480

[25] Douwe Korff, Opinion on the implications of the exclusion from new binding European instruments on the use of AI in military, national security and transnational law enforcement contexts, written at the request of the European Center for Not-for-Profit Law (ECNL), October 2022. See in particular section 2.2, *The risks [of AI] spelled out further*

and linked to the Charter in the EU Court of Justice PNR judgment, available at:

https://ecnl.org/sites/default/files/2022-10/ECNL%20Opinion%20AI%20national%20security.pdf

[26] In particular the 1969 UN Convention on the Elimination of Racial Discrimination (CERD), the 1979 UN Convention on the Elimination of Discrimination Against Women (CEDAW), and the 1981 UN General Assembly Declaration on the Elimination of All Forms of Intolerance and of Discrimination Based on Religion or Belief. Note that beyond these areas of discrimination, Article 21 CFR prohibits discrimination on "*any ground such as sex, race, colour, ethnic or social origin, genetic features, language, religion or belief, political or any other opinion, membership of a national minority, property, birth, disability, age or sexual orientation*".

Cf. also the extensive work and standard-setting in this area carried out by the Council of Europe, noted at:

https://www.coe.int/en/web/inclusion-and-antidiscrimination/about

[27] Douwe Korff, A (Not So Simple) Question: Must All Algorithmic Discrimination Be "Prevented" and "Eliminated" or Need It only be "Minimised"?, 29 September 2025, available at SSRN:

https://ssrn.com/abstract=7409523

[28] ICO, Guidance on AI and data protection, updated version of 15 March 2023, p. 18, available at:

https://ico.org.uk/media2/ga4lfb5d/guidance-on-ai-and-data-protection-all-2-0-38.pdf

[29] See note 24, above.

[30] Article 29 Working Party, Opinion 03/2013 on purpose limitation (WP203), adopted on 2 April 2013, p. 15, available at:

https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf

[31] *Idem*, pp. 15 – 16, emphasis added. The Opinion provides examples in an Annex, to which footnotes in the quoted passage refers.

[32] *Idem*, 3rd example in the Annex, on p. 51.

[33] Douwe Korff, *Applying the EU General Data Protection Regulation (GDPR) to the use of bank customer transaction data to train an AI based marketing model: A deeply defective and flawed decision of the Belgian data protection authority*, 11 April 2024, available at:

https://www.ianbrown.tech/wp-content/uploads/2024/04/KORFF-D-A-bad-Belgian-decision-240411-EDITED-1.pdf

Douwe Korff, *Applying the EU General Data Protection Regulation (GDPR) to the use of personal data to train an AI system: A follow-up to my note on a deeply defective and flawed decision of the Belgian data protection authority*, 16 April 2024, available at:

https://www.ianbrown.tech/wp-content/uploads/2024/04/KORFF-D-AI-the-GDPR-a-follow-up.pdf

[34] Note that **GPAI** is different from both "**Artificial General Intelligence**" (AGI) and "**Generative AI**" (**GenAI)**. These are defined by Stanford HAI as follows:

"***AGI*** *stands for* ***Artificial General Intelligence****, which means an AI system with general, human-level (or beyond) ability to learn, reason, and apply knowledge across a wide range of tasks and domains. AGI systems conceivably could handle novel situations, not just perform well on a single, narrow task. The term is controversial in several ways, including that different people mean different things by "human-level intelligence," and there's no universally accepted test, so claims are hard to verify. There are also safety and ethical concerns debated by AI experts.*" See:

https://hai.stanford.edu/ai-definitions/what-is-agi-artificial-general-intelligence

AGI is still essentially a theoretical type of AI.

"***Generative AI*** *(or* ***GenAI****) refers to AI systems that can create new content like text, images, music, code, or video. These systems learn patterns from training data and generate novel outputs that resemble the original data, often powered by architectures like GANs, transformers, diffusion models, and variational autoencoders. These models power applications including chatbots, code generation, and creative tools. They also raise questions about the potential for misuse including creating misinformation and deepfakes.*" See:

https://hai.stanford.edu/ai-definitions/what-is-generative-ai

On the differences between broad GPAI and narrow AI systems, see: ArcticSledge, What Is General Purpose AI: The Complete Guide to GPAI in 2026, available at:

https://www.articsledge.com/post/general-purpose-ai-gpai

[35] And in due AGI: see the definition in the previous note.

[36] See European Commission, *General-Purpose AI Models in the AI Act – Questions & Answers*, lastly updated on 20 November 2024, available at:

https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers

See also AI & Partners, EU AI Act: Classification of General-Purpose AI Models as General-Purpose AI Models with Systemic Risks – Methodology, December 2024, available at:

https://futurium.ec.europa.eu/system/files/2024-12/EU%20AI%20Act%20-%20Methodology%20-%202024-12-11.pdf

[37] Christopher Kuner (ed.) *et al*., The EU General Data Protection Regulation (GDPR: A commentary , 1st ed., OUP 2020, p. 1246.

[38] Christopher Kuner (ed.) *et al*., note 35, above), p. 1246.

[39] See Art. 89(4). [original footnote]

[40] *Idem*.

[41] EDPB, Guidelines 1/2026 on processing of personal data for scientific research purposes, adopted on 15 April 2026 (v.1 (for public consultation), available at:

https://www.edpb.europa.eu/system/files/2026-05/edpb_guidelines_202601_scientificresearch_en.pdf

[42] See section 2.1 of the guidelines, on pp. 8 – 10, where each of these factors is expanded upon, with references to authoritative international standards such as the Helsinki Principles on medical research and the European Commission's Ethics in Social Science and Humanities.

[43] Para. 12, on p. 10.

[44] Example 2 on p. 11.

[45] See note 33, above.

[46] Cf. Recital 159 that reads in the relevant part: "*For the purposes of this Regulation, the processing of personal data for scientific research purposes should be interpreted in a broad manner including for example technological development and demonstration, fundamental research, applied research and privately funded research*." In Germany, it has been suggested that the concept of "scientific research" can and should be more clearly (narrowly) defined by reference to the case-law of the Constitutional Court, which limits it to research by independent (research) institutions (German: *unabhängige und selbständige Institutionen*). See Gola/Heckmann, DS-GVO BDSG Kommentar, 3rd ed.,, Beck Verlag, 2022, Commentary on Art. 89 DS-GVO (=GDPR), margin number. 23, on p. 1068, with reference to other authors. But they agree that this is not a limitation that can be read into the GDPR.

[47] But note that the Article 6 "legitimate interest" legal basis cannot be relied upon by public authorities to process any personal data, let alone sensitive data (Article 6(1), final sentence).

[48] See sub-section *i.*, above. I discuss the further nuances added to this in the EDPB guidelines on anonymisation in section 3.2, below.

[49] In an Article 29 Working Party Advice Paper on special categories of data ("sensitive data") (Reference Ares(2011)444105), issued on 20 April 2011, that discussed various options to keep, amend or leave out a closed list, and the perceived advantages and disadvantages of the various options, the Working Party rightly noted that: "According to recent scientific thought, there are no races. In biological terms, Homo sapiens today is not divided into different races nor sub-types." (p. 10). Available at:

https://ec.europa.eu/justice/article-29/documentation/other-document/files/2011/2011_04_20_letter_artwp_mme_le_bail_directive_9546ec_annex1_en.pdf

This is reflected in recital 51 to the GDPR that stresses that "the use of the term 'racial origin' in this Regulation does not imply an acceptance by the Union of theories which attempt to determine the existence of separate human races."

[50] EDPB, Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, version 1.0 (for consultation), adopted on 8 October 2024, p. 14, emphases added.

[51] CJEU, judgment of 4 July 2023, Case C-252/21, *Meta v. Bundeskartellamt* (ECLI:EU:C:2023:537), para. 89. [original footnote]

[52] Ibid., para. 68. [original footnote]

[53] Ibid., para. 69. [original footnote]

[54] This was – and is – an issue in relation to the: USA interception of Internet communications in bulk, although that was glossed over in the drafting and adoption of the adequacy decision on that third country. It should also have been taken into account in relation to the UK now that it is also a third country in EU terms, because the intercepting of the Internet communications from the data carriers is in fact done by the USA and the UK working together. See: Ian Brown & Douwe Korff, Exchanges of personal data after the Schrems II judgment, study carried out at the request of the European Parliament's Civil Liberties (LIBE) Committee into the future of EU – US flows of personal data, July 2021, available at:

https://www.europarl.europa.eu/RegData/etudes/STUD/2021/694678/IPOL_STU(2021)694678_EN.pdf

[55] Cf. the concerns expressed in the USA over the selling of location data on people seeking reproductive and gender healthcare to anti-abortion groups. See the Electronic Frontier Foundation blog by Karen Gullo, *Location Data*

Tracks Abortion Clinic Visits. Here's What to Know, 15 March 2024, available at:

https://www.eff.org/deeplinks/2024/03/location-data-tracks-abortion-clinic-visits-heres-what-know

[56] Article 29 Working Party, Advice Paper on special categories of data ("sensitive data") (footnote 298, above), p. 6, original italics, emphasis in bold added.

[57] *Idem*, p. 8, in the section discussing various "problems … in applying the concept of sensitive data".

[58] See: CJEU, judgment of 4 October 2024 in Case 446/21, *Maximilian Schrems v. Meta Platforms Ireland Ltd, formerly Facebook Ireland Ltd* (ECLI:EU:C:2024:834), in particular paras. 81 - 83.

[59] See the EDPB Guidelines 1/2026 on processing of personal data for scientific research purposes (footnote 40, above), section 4.1.2.1, *Broad consent*.

[60] This section again draws on my earlier paper, Legal Requirements of Transparency, Explainability and Procedural Fairness in the Taking of Automated/AI-based Decisions that Affect Individuals (note 21, above).

[61] Cf. the reference in para. 14 of the *SCHUFA* judgment (referenced in note 65, below), quoted later in the text, to "*[a credit scoring agency's algorithm] establish[ing] a prognosis on the probability of a future behaviour of a person ('score') … based on certain characteristics of that person, on the basis of mathematical and statistical procedures [and] the assumption that, by assigning a person to a group of other persons with comparable characteristics who have behaved in a certain way, similar behaviour can be predicted.*"

[62] See notes 23, 24 and 25, above.

[63] CJEU, First Chamber judgment of 7 December 2023 in Case 634/21, *OQ v. Land Hessen & SCHUFA Holding AG*, ECLI:EU:C:2023:957 ("*SCHUFA* judgment"), para. 73, repeated verbatim in the conclusion in para. 75. See paras. 40 – 72 for the detailed reasoning underpinning this conclusion.

[64] Douwe Korff, Legal Requirements of Transparency, Explainability and Procedural Fairness in the Taking of Automated/AI-based Decisions that Affect Individuals (note 21, above), section 4.6, *The level of detail to be revealed* (pp. 16 – 18).

[65] Berliner Beauftragte für Datenschutz und Informationsfreiheit (BlnBDI) (Berlin DPA), *Computer sagt nein* (*Computer says no*), press release, 31 May 2023, available at:

https://www.datenschutz-berlin.de/pressemitteilung/computer-sagt-nein/

[66] Cf. again the detailed discussion of when a data subject's consent can be valid in the WP29 Guidelines on consent under [the GDPR] (footnote 59, above).

[67] See the postscript to my paper, A (Not So Simple) Question: Must All Algorithmic Discrimination Be "Prevented" and "Eliminated" or Need It only be "Minimised"? (note 27, above), pp. 10 – 13. This referenced Graham Smith, *Ofcom's proactive technology measures: principles-based or vague?*, 4 August 2025, avilable at: https://www.cyberleagle.com/2025/08/ofcoms-proactive-technology-measures.html

[68] This relates to the following Ofcom consultation:

https://www.ofcom.org.uk/online-safety/illegal-and-harmful-content/online-safety-additional-safety-measures

[69] This is a composite definition, based on the definitions of "*racial discrimination*", "*discrimination against women*" and "*intolerance and discrimination based on religion or belief*" in the 1969 UN Convention on the Elimination of Racial Discrimination (Article 1(1)), the 1979 UN Convention on the Elimination of Discrimination Against Women (Article 1), and the 1981 UN General Assembly Declaration on the Elimination of All Forms of Intolerance and of Discrimination Based on Religion or Belief (Article 2(2). The list of protected characteristics in square brackets is taken from the non-discrimination article in the EU Charter of Fundamental Rights, Article 21.

[70] See again my paper, A (Not So Simple) Question: Must All Algorithmic Discrimination Be "Prevented" and "Eliminated" or Need It only be "Minimised"? (note 27, above).

[71] Cf. the full definition in Article 2(2) GDPR:

"'***processing'*** *means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage,* ***adaptation or alteration****, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction.*"

[72] 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 (WP248rev0.1), adopted on 4 April 2017, as last revised and adopted on 4 October 2017, available at:

https://ec.europa.eu/newsroom/article29/items/611236/en

They were endorsed by the EDPB in Endorsement 1/2018 (note 60, above).

[73] 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**), OJ L, 2024/1689, 12.7.2024, ELI:

http://data.europa.eu/eli/reg/2024/1689/oj

[74] See: European Commission, *Timeline for the Implementation of the EU AI Act*, available at:

https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act

The coming into application dates were amended and extended by the Digital Omnibus on AI, which is the part of the overall Digital Omnibus package that has already been adopted in the form of Regulation (EU) 2026/1744. I discuss the still under consideration parts of the DO in section 3, below.

[75] European Commission, *Shaping Europe's digital future – AI Act*, last updated on 3 August 2026, at:

https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

For another broad overview, see the Future of Life Institute's *High-level summary of the AI Act*, 27 February 2024, updated on 30 May 2024 in accordance with the Corrigendum version of the AI Act published in the OJ on 12 July of that year, available at:

https://artificialintelligenceact.eu/high-level-summary/

[76] European Commission, *AI Act Single Information Platform*, available at:

https://ai-act-service-desk.ec.europa.eu/en

This provides links to pages with an *AI Act Explorer*, an AI Act *Compliance Checker*, the Commission's *AI Act Service Desk*, *Latest FAQs* and *Latest News*, etc.

[77] Note that the Act makes a distinction between "AI *systems*" and "AI *models*". I discuss the intention and implications of this at xxx, below.

[78] European Commission, *Shaping Europe's digital future – AI Act* (note 78, above).

[79] The Commission published two key documents to support the practical application of the prohibited practices:

- The guidelines on prohibited AI practices under the AI Act, which offer legal explanations and practical examples to help stakeholders understand and comply with the prohibitions.
- The guidelines on the AI system definition of the AI Act, to assist stakeholders in determining the scope of the AI Act.

I discuss the latter briefly in sub-section *ii*.

[80] Note that "biometric identification" (remote or otherwise) is distinguished from "biometric verification" or "-authentication". As clarified in Article 3, paras. 35 and 36:

*(35) '**biometric identification**' means the automated recognition of physical, physiological, behavioural, or psychological human features for the purpose of establishing the identity of a natural person by comparing biometric data of that individual to biometric data of individuals stored in a database;*

*(36) '**biometric verification**' means the automated, one-to-one verification, including authentication, of the identity of natural persons by comparing their biometric data to previously provided biometric data.*

[81] Future of Life Institute, *High-level summary of the AI Act*, updated on 31 August 2026 in accordance with the amendments to the AI Act adopted as part of the Digital Omnibus on AI, available at:

https://artificialintelligenceact.eu/high-level-summary/

[82] European Commission, *Shaping Europe's digital future – AI Act* (note 75, above).

[83] See note 77, above, and again the discussion at xxx, below.

[84] The same was the case with the 1995 EC Data Protection Directive. One of the reasons for replacing it with the GDPR was precisely to put fundamental rights protection foremost, and the good functioning of the internal market second.

[85] See note 84, above.

[86] Regulation (EU) 2019/1020 of the European Parliament and of the Council of 20 June 2019 on market surveillance and compliance of products and amending Directive 2004/42/EC and Regulations (EC) No 765/2008 and (EU) No 305/2011, OJ L 169, 25.6.2019, pp. 1–44, ELI: http://data.europa.eu/eli/reg/2019/1020/oj

[87] See Harris, O'Boyle & Warbrick, Law of the European Convention on Human Rights, OUP, 2nd ed. (2009), Chapter 8, section 3.1, *'In accordance with the law'/'prescribed by law'*. (This general section was not retained in the subsequent 3rd edition of 2014.) I have also summarised the requirements in a simple student handout: Douwe Korff, The standard approach under articles 8 -11 ECHR and article 2 ECHR, 1 August 2008, available at SSRN: https://ssrn.com/abstract=5038559

[88] European Commission, *Shaping Europe's digital future – AI Act* (note 78, above).

[89] Opinion on the implications of the exclusion from new binding European instruments on the use of AI in military, national security and transnational law enforcement contexts, written at the request of the European Center for Not-for-Profit Law (ECNL) (hereafter: "**ECNL Opinion**"), October 2022, available at:

https://ecnl.org/sites/default/files/2022-10/ECNL%20Opinion%20AI%20national%20security.pdf (full opinion)

https://ecnl.org/sites/default/files/2022-10/ECNL%20Opinion%20AI%20national%20security_exec%20summary_0.pdf (executive summary)

[90] See, e.g. :

European Parliament, Collection of studies prepared by Policy Department C for the PEGA Committee: *Spyware as a threat to fundamental rights in the EU*, , April 2024, available at:

https://www.europarl.europa.eu/RegData/etudes/BRIE/2024/761472/IPOL_BRI(2024)761472_EN.pdf

No place for Palantir in a democratic Europe!, Joint call by over 120 campaigning organisations and unions across Europe to governments and the European Commission to scrap all existing public sector contracts with the US tech company Palantir, 28 September 2026, full letter available at:

https://www.globaljustice.org.uk/wp-content/uploads/2026/09/LETTER_SIGNATORIES.pdf

[91] ECNL Opinion (note 89, above), p. 27, words in square brackets added This conclusion in turn rests on a study by Ian Brown and me for the European Parliament, Exchanges of Personal Data After the Schrems II Judgment, 12 July 2021, section 2.2.2, *The national security exemption in the EU Treaties* (pp. 22 – 24), available at SSRN:

https://ssrn.com/abstract=3884896

[92] I am quoting the third point from this EDAP factsheet:

https://www.eeas.europa.eu/sites/default/files/defence_package_factsheet_0.pdf

[93] See the study for the European Parliament, Isabelle Ioannides (ed.), EU Defence Package: Defence Procurement and Intra-Community Transfers Directives, October 2020, available at:

https://op.europa.eu/en/publication-detail/-/publication/3a977249-3e88-11eb-b27b-01aa75ed71a1/language-en

[94] ECNL Opinion (note 89, above), pp. 36 – 37.

[95] See, e.g., Marco Sassòli, Legitimate Targets Of Attacks Under International Humanitarian Law, Cambridge, 2003, available at:

https://hhi.harvard.edu/files/humanitarianinitiative/files/session1_legitimate_targets_ihl.pdf?m=1615827575

(original footnote)

[96] BBVA Open Mind, Drones That Kill on Their Own: Will Artificial Intelligence Reach the Battlefield?, 8 May 2018, available at: https://www.bbvaopenmind.com/en/technology/artificial-intelligence/drones-that-kill-on-their-own-will-artificial-intelligence-reach-the-battlefield/ (original footnote)

[97] Call for tenders for Automatic Target Recognition of Personnel and Vehicles from an Unmanned Aerial System Using Learning Algorithms, 29 November 2017, available at:

https://www.sbir.gov/sbirsearch/detail/1413823 (original footnote)

The above page (on an official US Government website) is no longer up, but of course such research is continuing.

[98] https://www.stopkillerrobots.org/ (original footnote)

[99] 972 Magazine, 3 April 2024, available at:

https://www.972mag.com/lavender-ai-israeli-army-gaza/

[100] ECNL Opinion (note 89, above), p. 36, words in square brackets added.

[101] *Idem*, p. 32. Words in square brackets added.

[102] In section 2.1.*iii*, above, I wrote that in my opinion GPAI models or systems are inherently incompatible with the purpose-specification and -limitation principle in the GDPR, but this view is clearly not shared by the drafters of the AI Act. The further exemption in relation to "pre-market" GPAI models aggravates this issue.

[103] CNN, *An OpenAI test model escaped and broke into a real company's servers*, 22 July 2026, available at: https://www.cnn.com/2026/07/22/tech/openai-hugging-face-ai-cybersecurity

[104] TechInsider, *Google Confirms 3 AI Escapes as NYC Grills Big Tech*, 7 October 2026, available at:

https://tech-insider.org/google-confirms-3-ai-agent-escapes-nyc-hearing-2026/

[105] See, e.g.:

Kai Nicol-Schwar, *These four charts show how reliant Europe is on U.S. digital infrastructure*, CNBC, 1 February 2026, available at:

https://www.cnbc.com/2026/02/13/four-charts-europes-reliance-us-digital-infrastructure.html?msockid=2bcb5f0f4faf6c312b8348764eed6d2a

The Conversation, *Europe wants to end its dangerous reliance on US internet technology*, 21 January 2026, available at:

https://theconversation.com/europe-wants-to-end-its-dangerous-reliance-on-us-internet-technology-274042

Neil C Hughes, *Europe's risky reliance on US tech*, CyberNews, 12 February 2026, available at:

https://cybernews.com/security/europes-risky-reliance-us-tech-digital-shutdown/

[106] Cf. the discussion of *CJEU requirements relating to indirect or direct access to personal data by third country intelligence agencies* in the study by Ian Brown and me for the European Parliament on Exchanges of personal data after the Schrems II judgment (note 54, above), pp. 41 – 45.

[107] As noted in section 2.1.*ii*, above, with reference to my paper asking the question A (Not So Simple) Question: Must All Algorithmic Discrimination Be "Prevented" and "Eliminated" or Need It only be "Minimised"? (note 27, above), that under international human rights law, the Charter and the GDPR (and its sister instruments) discrimination must not just be “minimised” but actually “elmininated”; and that de-biasing should therefore also not stop at “minimising” or “mitigating” biased outputs, but should actually correct them to ensure they do not lead to biased outcomes. The reference to “mitigating” of biases in Article 10 again ignores this important issue.

[108] European Commission, *Guidelines on transparency obligations for providers and deployers of certain AI systems*, last updated on 6 August 2026, available at:

https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-transparency-obligations

The actual guidelines and associated information are available at:

https://digital-strategy.ec.europa.eu/en/library/guidelines-transparency-obligations-providers-and-deployers-ai-systems

[109] See notes 6 – 9, above.

[110] See: GDPRbuzz, *EU Council Rejects Changes to GDPR Personal Data Definition*, 26 February 2026, available at: https://gdprbuzz.com/news/eu-council-rejects-changes-to-gdpr-personal-data-definition/

[111] CJEU, *SRB* judgment (note 19, above), para. 86, quoted on p. 4, above.

[112] Douwe Korff, Digital Omnibus: Proposed Changes to the Definition of Personal Data in the EU GDPR Analysed in the Light of the SRB and Earlier CJEU Judgments, 9 December 2025, available at SSRN:

https://ssrn.com/abstract=5891404

Douwe Korff, The Concepts of "Data Subject", "Personal Data", "Pseudonymous Data" and "Anonymous Data" in the EU GDPR, 11 November 2025, available at SSRN: https://ssrn.com/abstract=5734864

[113] EDPB, Guidelines 02/2026 on Anonymisation, v.1 for consultation, available at:

https://www.edpb.europa.eu/system/files/2026-07/edpb_guidelines_202602_anonymisation_v1_en_0.pdf

The consultation is open until 30 October 2026, see:

https://www.edpb.europa.eu/public-consultations/guidelines-022026-on-anonymisation_en#no-back

[114] EDPB, Guidelines 02/2026 on Anonymisation (previous note), section 2.5.2, paras. 37 – 42, pp. 15 – 16.

[115] See, e.g., C-413/23 P *EDPS v SRB* and the effects of anonymisation (or lack thereof) on certain aspects of the information obligation as discussed in Example 2. (original footnote)

[116] See also EDPB Guidelines 1/2026 on processing of personal data for scientific research purposes, adopted on 15 April 2026, version for public consultation, section 3.1. (original footnote)

[117] For more information on the legitimate interest balancing test, see EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR. (original footnote)

[118] For further guidance on this topic, see EDPB Guidelines 9/2022 on personal data breach notification under GDPR. (original footnote)

[119] I have reproduced here the most salient paragraphs, with minor edits. For the full discussion in the guidelines, see section 3.4, *The three criteria*, paras. 52 – 82, on pp. 18 – 25, with several examples.

[120] See the entries “infer” and “inference” in Catherine Soanes & Angus Stevenson (eds.) Concise Oxford English Dictionary (11th edition Revised, OUP 2009). (original footnote)

[121] In this regard, it is important to recall that, as discussed in sub-sections 2.3 and 2.4, information may be considered personal data, even if its link to an individual is not immediately apparent or if the information does not itself contain everything necessary to identify the individual to whom it relates. (original footnote)

[122] Here, “information about the population at large” means any large group of people about whom information is not collected individually, but extrapolated from a sample, with any connection to the individuals left behind. In technical terms, such inferences can be seen as part of the “prior belief”. (original footnote)

[123] In more technical terms, the argument presented in this paragraph is as follows: Inferring means making a prediction on a property S having observed the given data A(D) and additional information, where A(D) was produced by running the anonymisation process A to the original data D and the additional information serves as a conditioning factor. Suppose that S is a property about a person named John Smith and that data relating to that person is included in D. We denote by D' a copy of the dataset D, with the only difference that all data relating to John Smith have been removed. Therefore, D' does not contain any data relating to John Smith. It may still be possible to make some inference on John Smith based on A(D'), including on property S. For example, this may be possible if the other records in D' relate to persons that are "similar" to John Smith (e.g. because they live in the same city). However, this inference cannot depend on John Smith's data, as there is no data relating to John Smith in D'. Therefore, if the inference on John Smith works in the same way for A(D) and A(D'), it is not an inference drawn from his data but from data relating to the population at large, for which D' is representative. (original footnote)

[124] In such a case, "Prior belief" (…) could include knowledge coming from an alternative version of the given data which has been produced without the record of the persons that are being identified – A(D') in [the previous footnote] – but which still applies to them and other persons by way of extrapolation. Inferences from the given data should yield a significantly different "posterior belief" (on which an action could turn) to be considered related to the individual in a way that classifies the given data as personal. (original footnote)

[125] Example 14, given in para. 71, on p. 23.

[126] Cf. para. 99 in the guidelines.

[127] See also EDPB: Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models. (original footnote)

[128] Simon Lermen, Daniel Paleka, Joshua Swanson, Michael Aerni, Nicholas Carlini and Florian Tramèr, Large-scale online deanonymization with LLMs, 25 February 2026, available at:

https://arxiv.org/html/2602.16800v2

[129] HERALD, *LLMs Just Killed Your Online Anonymity—And It's Terrifyingly Cheap*, 3 March 2026, available at:

https://chyshkala.com/blog/llms-just-killed-your-online-anonymity-and-it-s-terrifyingly-cheap?shem=aimgspe

[130] Arvind Narayanan & Vitaly Shmatikov, "Robust De-anonymization of Large Sparse Datasets" (2008) 2008 IEEE Symposium on Security and Privacy 111, https://doi.org/10.1109/SP.2008.33 (referenced in footnote 39 in the guidelines).

[131] As explained in section 1, the Digital Omnibus Regulation is separate from the Digital Omnibus on AI that was adopted in the form of Regulation (EU) 2026/1744: see note 74, above.

[132] See the "Legislative Train Schedule" on the Digital Omnibus Regulation Proposal, available at:

https://www.europarl.europa.eu/legislative-train/theme-a-new-plan-for-europe-s-sustainable-prosperity-and-competitiveness/file-digital-package

- o – O – o -