The EDPB's Rorschach Test: What the data protection body's Opinion on AI training Means for GDPR Enforcement

Author: Itxaso Domínguez de Olazábal, PhD, Policy Advisor, EDRi

Why has EDPB's long awaited Opinion on AI training sparked intense debate?

On 4 September 2024, the Irish data protection watchdog invoked Article 64(2) of the GDPR to request an examination from the European Data Protection Board (EDPB)—an umbrella body that ensures consistent application of the GDPR by coordinating the work of all EU/EEA watchdogs—**regarding the processing of personal data in the context of AI training**. This followed a number of complaints lodged by EDRi member noyb, which successfully challenged the practices that Meta, X and others had started implementing in the EU to feed their AI databases.

On 17 December 2024, the EDPB released its long-awaited 'Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models.' **The process has highlighted some insufficiencies: an expedited timeline, limited and asymmetrical stakeholder involvement, and a problematic framing**. In terms of content, this document has sparked intense debate, and it will continue to do so in the near future. The main reason for this is that, **the more you reflect on it, the more it resembles a Rorschach test—everyone seems to see what they want to see**.

At first glance, the Opinion might appear to offer a strong reaffirmation of GDPR principles, countering significant industry pressure to relax safeguards. It seems to position data protection rights at the centre of the ongoing AI conversation, advocating for robust regulation. But **many digital rights advocates were concerned because the Opinion focused on 'Legitimate Interest' (LI) as a basis for processing data, which weakens the requirement for companies to get explicit consent from people.** This **opened a door for companies to use personal data without permission**, making it easier to train AI systems on people's information, allowing them to bypass privacy safeguards, leading to potential misuse of data, surveillance, and bias in AI (not to mention the environmental costs!). In the long run, this wouldn't just weaken the foundations of the data protection framework but could also have **far-reaching consequences for other fundamental rights**. The personal data in question fuels the algorithms that big tech social media platforms use to exploit our behaviours, shape perceptions, and manipulate vulnerabilities.

Legitimate Interest: A Loophole in the Making?

A central debate in AI regulation today revolves around the use of LI as a legal basis for processing personal data. While consent is another legal basis, it is often dismissed by companies as inconvenient or impractical—precisely because it is a more rigid requirement. Consent needs to meet specific criteria set out by the regulation, including being freely given, specific, informed, and unambiguous. **Businesses have long sought to avoid or manipulate consent mechanisms, favouring alternatives that grant them more flexibility**. In this regard, LI is perceived as more flexible, allowing organisations to process personal data based on their interests, provided these interests are not overridden by the rights and freedoms of individuals. **Although the GDPR does not explicitly**

favour one legal basis over another, prioritising consent aligns more closely with a rights-based approach, ensuring that data protection is not sacrificed for the convenience of data-driven industries. And the EDPB has rightly warned several times against treating LI as a catch-all justification for data processing.

This debate is particularly significant because **LI has been framed as a compromise—one that sidesteps reopening the GDPR while ostensibly making AI development more 'workable'** in the EU, particularly in response to industry demands for regulatory flexibility. The Opinion played along with this framing by omitting any mention of consent, yet it also took a subtly strategic approach: **while allowing for the possibility of bypassing consent, it set a high bar for doing so**. It underscored that any reliance on LI must strictly adhere to GDPR criteria, with a particular emphasis on data minimisation—an essential safeguard in the AI context—and must not come at the expense of individuals' rights.

The digital rights community has long warned that, if misused, LI can serve as a loophole that erodes transparency and accountability, weakening GDPR's protective framework. A key concern is that **companies may increasingly rely on 'mitigating measures'—meant to minimise the impact of data processing on individuals—as a way to justify non-compliance with GDPR principles**. While such measures are important in reducing harm, they must not become a convenient tool for industry players to claim compliance while continuing harmful data practices. If mitigating measures are used merely to legitimise LI-based processing, they risk undermining the GDPR's fundamental protections. Instead, **regulators must ensure that companies adhere fully to data protection principles and rely on alternative legal bases where LI fails to meet the strict standards set by the GDPR**. Anything less would enable the continued erosion of individuals' rights under the guise of regulatory flexibility.

While industry players have welcomed the Opinion's treatment of LI, history has shown that such acceptance often comes with a hidden agenda. This has been a **persistent issue with GDPR enforcement: for the regulation to be meaningful, its rules must be applied fully and consistently**. Yet, in practice, enforcement remains weak, and even minor loopholes are exploited. In an ecosystem dominated by companies with insatiable data appetites, **rights are routinely sidelined in favour of superficial compliance measures** that create the illusion of adherence to the law rather than ensuring substantive data protection.

The Opinion also reminds us that LLMs, despite not functioning like traditional databases (more on this below), are still subject to the same core data protection principles. It's crucial to highlight that the argument some in the industry make—that LLMs are not databases and therefore do not need to adhere to the same data protection rules—does not hold up. The Opinion reinforces that **the absence of a traditional data storage system in LLMs does not absolve them from compliance with GDPR requirements**. Instead, it highlights the importance of tailored, robust mitigation measures that align with the nature of AI technologies. The issue should not be framed as an excuse to weaken compliance, but rather as a challenge to develop more effective safeguards that fit the complexities of AI.

The Myth of Data Anonymisation in AI: A Cautionary Tale

The DPC had asked the EDPB to answer the following question: 'Is the final AI Model, which has been trained using personal data, in all cases, considered not to meet the definition of personal data (as set out in Article 4(1) GDPR)?' There were months of heated debates about this on LinkedIn and other spaces where data protection enthusiasts converge—debates notably sparked by a discussion paper from the Hamburg Data Protection Authority (DPA), which argued that Large Language Models (LLMs) do not store personal data and therefore the mere storage of an LLM does not constitute processing. Now, the EDPB's Opinion has clarified an important point. It rightly recognised that **AI models trained on personal data cannot automatically be considered anonymous.** This is because such models may memorise and reproduce elements of the data they were trained on, raising concerns about potential re-identification and the processing of personal data.

This clarification was crucial because it **reaffirmed that transforming personal data through AI training does not eliminate GDPR obligations**. If AI models can retain, reproduce, or enable the re-identification of personal data, **they remain subject to data protection rules**. This recognition is essential to prevent companies from using AI as a loophole to evade legal responsibilities, ensuring that individuals' rights over their data are upheld—even in contexts that industry often portrays as too technically complex for regulation to apply.

The EDPB seems to indicate that anonymisation is not outright impossible simply due to the presence of residual risks. Rather, it highlights the need for data controllers to perform an additional, detailed evaluation of these risks to determine whether they can be effectively managed, further emphasising that anonymisation processes must comply with the principle of accountability whereby developers should document the technical and organisational measures taken to minimise identification risks. However, **the Opinion should have explicitly stated that achieving true anonymisation of data is not just challenging—it is, in fact, virtually impossible for any of the current commercial model families, without worsening their performance considerably**. While this would have been the most accurate stance, it likely wouldn't have been politically acceptable in the current climate (more on this below). After all, this is merely an 'Opinion', but any DPA diverging from it would need to justify why.

This recognition serves as an important **reminder to approach controllers' claims of anonymity with caution**. Too often, such assertions are used as a veneer to sidestep compliance, even when the underlying practices fail to meet the necessary standards. The EDPB's stance is welcome in this regard, reminding us that the functionality of AI should not come at the expense of fundamental rights. This perspective is especially needed to **counter the misleading argument promoted by industry and others, which asserts that AI models are not databases of structured facts or personal data, and that they do not work by pulling data from a database or by 'copying and pasting' existing content**.

The EDPB Opinion makes it clear that **the lawfulness of AI training cannot be assessed in isolation from its subsequent deployment.** If an AI model was developed using unlawfully processed personal data, this illegality does not simply disappear at later stages. This is particularly relevant in the case of DeepSeek, which is reportedly using personal data originally illegitimately acquired by companies like OpenAI. Even **if DeepSeek were to claim that the data was anonymised at a later stage, this would not absolve it—or any controller deploying the model—from the obligation to assess whether the AI system was built on GDPR violations**. The Opinion explicitly warns that

controllers must take into account whether the AI model is the result of an infringement of the GDPR, particularly if such a violation has been recognised by a supervisory authority or court. The fact that the data was unlawfully obtained in the first place should weigh heavily in assessing the lawfulness of its subsequent use.

This is closely tied to a broader issue raised by the Opinion: the **problematic assumption that later anonymisation can retroactively cleanse unlawful data processing**. In theory, one could argue that if personal data are anonymised after being unlawfully processed during development, the subsequent deployment phase would no longer be tainted by the initial illegality. However, this approach presents significant practical and legal challenges. Anonymisation must be assessed in context, meaning that whether data can truly be considered anonymised depends on factors such as the risk of re-identification, how the data are used, and the broader environment in which they are processed. Moreover, AI systems rarely follow a linear or predictable path, making it unclear whether anonymisation at a later stage would truly neutralise the effects of earlier unlawful actions. **The Opinion reinforces that controllers cannot rely on anonymisation as a loophole to evade accountability, particularly if they oversaw both the unlawful data collection and the subsequent use of the AI model.** Failing to impose consequences for earlier violations would create perverse incentives, encouraging companies to disregard data protection principles in training, knowing they could later claim compliance through anonymisation.

All of this leads us to an important point about excessive reliance on *post-facto* anonymisation. **Over-relying on anonymisation as a safeguard, an implicit recommendation of the Opinion, could lead to complacency and further legitimising the existing check-box compliance culture.** Instead of relying on post-facto anonymisation, it is crucial for developers to focus on designing AI systems that adhere to data minimisation principles from the outset, ensuring that personal data are collected and used only when absolutely necessary. **The digital rights community has long argued that the focus on data protection in AI is not about stifling competitiveness, but rather about creating space for privacy-preserving alternatives to thrive.** Far from hindering innovation, prioritising privacy in AI development can drive the creation of new technologies which allow AI systems to function without relying on vast, unregulated data pools. This shift **challenges the prevailing belief that more data equals better AI and opens the door for a more sustainable, ethical AI landscape** that fosters trust and meets regulatory standards. By encouraging solutions that respect fundamental rights, we can cultivate a competitive market where privacy and innovation go hand in hand, ensuring that technological progress does not come at the cost of individuals' privacy.

What the Opinion reminds us of, and importantly so (even if not explicit), is that **LLMs can indeed be developed without relying on data sourced from the internet.** This is crucial because a flawed argument gaining traction these days suggests that access to quality data is indispensable for improving AI outputs, mitigating social biases, and reflecting the diversity of societies. While this narrative seems compelling on the surface, it risks normalising invasive data practices under the guise of progress. **The diversity and accuracy of AI systems should not come at the cost of exploiting individuals' personal data without sufficient safeguards.** Equating the quality of AI with unrestricted access to data ultimately undermines both trust and the very protections designed to safeguard individuals.

DPAs and the Uneven Application of GDPR: The Price of Excessive Leeway

The one aspect we can all agree on is that **the Opinion grants excessive discretion to DPAs**. While the text acknowledges that all parties would 'greatly benefit from reaching a common position on the matters raised by this Request,' the excessive leeway given to DPAs risks leading to fragmented enforcement and inconsistent interpretations across jurisdictions. **This lack of uniformity has already proven to be a challenge under the GDPR since its introduction, undermining its effectiveness in safeguarding fundamental rights.** Clear, harmonised standards are essential to ensure that digital rights are upheld consistently across the EU. Unfortunately, this is an area where the Opinion falls short. That doesn't mean the Opinion is entirely unworkable, but its true value will depend on how it is interpreted and applied by the relevant regulators.

Why is this problematic? We are already witnessing a cycle of **reactive, piecemeal regulation—effectively a game of compliance whack-a-mole—where issues are addressed individually as they arise, without a clear, unified strategy**. This approach leaves critical gaps in enforcement, enabling companies to exploit regulatory uncertainty and continue processing personal data in ways that undermine individuals' rights. Take, for example, the varying interpretations by DPAs on data protection obligations in AI training. Italy's data protection authority fined OpenAI, the owner of ChatGPT, for processing personal data without a valid legal basis. On the other hand, the French DPA's position on AI training seems to favour industry interests. These inconsistencies can lead to confusion and, worse, a lack of accountability.

To truly understand the implications of this, we need to consider the political context surrounding the EDPB. As Lisette Mustert and Cristiana Santos have expertly pointed out, **the EDPB is supposed to operate with full independence in carrying out its responsibilities**. However, it does not always take a proactive approach in issuing clear guidance on specific issues, and the guidance that is released often fails to achieve the intended outcomes. As a result, **disagreements over the interpretation of the GDPR often persist, and its enforcement can appear inconsistent**. While we know the EDPB is unlikely to ban LLMs, one might wonder: who will take that responsibility, if anyone?

Finally, some critics argue that the EDPB Opinion lacks clear practical guidance for controllers. This critique misses an important point: **the process leading to the adoption of this Opinion has highlighted that some controllers have consistently failed, or chosen not, to demonstrate full compliance with GDPR obligations**. Rather than embedding GDPR principles into their operations, these controllers often shift the responsibility to regulators, expecting detailed implementation roadmaps or, more frequently, leniency in enforcement. The Opinion's reaffirmation of established GDPR standards serves as a reminder that these obligations are not new or unclear—they have been in place for years, and it is high time that they be fully integrated into AI operations.

Beyond Compliance: The Stakes of AI Regulation and the EDPB's Role

The broader narrative surrounding AI and regulation requires careful and nuanced examination, particularly in the context of the EDPB Opinion. **The persistent argument in the 'EU Bubble' that innovation and strong regulation cannot coexist is not only misleading but also potentially harmful.** A key counterpoint to the Draghi report's focus on competitiveness as a justification for easing GDPR protections—especially in relation to AI—is the reality that **sovereign-sized corporations now often wield power comparable to that of EU member states**. These entities can

shape regulatory processes and public policies in ways that undermine democratic governance. This power imbalance risks distorting the role of public institutions, steering them away from their essential duty to serve the public good and protect individuals' rights, and instead aligning them with private, profit-driven interests. **The assumption that AI developed by private companies will inherently serve the public good is, quite frankly, flawed**.

Diluting GDPR protections under the pretence of fostering innovation or enhancing competitiveness poses a serious threat to the democratic accountability of institutions. Such a shift could lead to a governance model in which private entities—rather than elected representatives—dictate the rules. This is not just an issue of data protection; it concerns the very foundations of democracy. Robust data protection standards are not merely about safeguarding individual privacy; they are about ensuring that public institutions remain independent and accountable, rather than becoming tools for private corporations pursuing profit over societal well-being. Yes, the Opinion, despite some insufficiencies, reinforces the idea that data protection should not be sacrificed at the altar of innovation, suggesting that responsible innovation can—and must—coexist with privacy protections. Yes, it could also be seen as reinforcing the GDPR as a cornerstone of a rights-respecting digital future. However, **its real impact will depend on how DPAs interpret it and, ultimately, how the courts rule**. This is where its **main—and most critical—flaw lies: given the persistent failure of DPAs to investigate complaints in a timely manner (or at all), individuals should not have to wait years for a court to reaffirm their fundamental rights.**

And this is because **the appearance of compliance that many controllers have adopted since the GDPR's implementation has too often been enough to satisfy these DPAs as regulators.** The risks and harms of this approach are not hypothetical—they are very real, ongoing, and extend far beyond the realm of data protection, especially in the context of AI. With six years of GDPR enforcement experience behind us, **there is a legitimate concern that the Opinion, while pushing back against some industry narratives, may leave too much room for exploitation and fail to provide adequate safeguards to ensure GDPR compliance in practice**. This highlights the critical need for civil society and other stakeholders to remain vigilant and advocate for stronger, harmonised enforcement. The stakes are simply too high to do otherwise.