Agnieszka Grzelak*

Reconciling Data Minimization with Model Maximization: Regulatory and Ethical Tensions in AI Development**

The rapid advancement of large-scale artificial intelligence (AI) systems, particularly large language models (LLMs), has created profound regulatory tensions in the realm of data protection. Central to this discourse is the conflict between the principles of data minimization, as enshrined in the General Data Protection Regulation (GDPR), and the data-intensive logic underpinning AI model development. This article explores some aspects of the legal, practical, and ethical implications of this tension from the perspective of data protection authorities (DPAs), analyzing current enforcement trends, regulatory guidance, and the prospective impact of the EU Artificial Intelligence Act. It argues that DPAs must evolve beyond traditional enforcement roles to become ethical stewards and proactive coordinators of AI governance in Europe to ensure that the fundamental principles are not weakened or ignored in the name of innovation.

***Keywords:** Data Minimization, AI Act, Large Language Models (LLMs), General Data Protection Regulation (GDPR), Ethical AI Governance.*

1. Introduction: Data Protection in the Age of Expansive AI Models

The increasing integration of artificial intelligence (AI) technologies into public services, private enterprise, and everyday life has intensified longstanding tensions between innovation and the protection of fundamental rights. Among the most acute

[*] Dr Habil of Legal Sciences, Professor at the Kozminski University in Warsaw (PhD 2000, The Jagiellonian University Cracow; Habil. 2016 Polish Academy of Science). Deputy President of the Personal Data Protection Office, Warsaw – Poland; a_grzelak@uodo.gov.pl. https://orcid.org/0000-0002-5867-8135

[**] The paper is the text of a keynote speech presented at the 33rd European Conference of Personal Data Protection Authorities ("Spring Conference"), hosted by the Personal Data Protection Service and held in Batumi. The information and views set out in this article are those of the author and do not necessarily reflect the official opinion of the European Commission.

of these is the emerging conflict between the GDPR’s foundational principle of data minimization and the data-intensive logic that drives the development of contemporary AI systems, especially large language models (LLMs).

The GDPR establishes data minimization as a bedrock principle of lawful data processing[1]. According to Article 5(1) (c), personal data must be "adequate, relevant and limited to what is necessary" in relation to the purposes for which it is processed. This mandate aims to curtail the collection and use of unnecessary personal data and to promote accountability, transparency, and respect for data subjects’ rights. However, in the context of AI model development, particularly LLMs, this principle is increasingly under strain.

LLMs function on the premise of scale: the broader and more diverse the dataset, the more nuanced and powerful the model becomes[2]. Developers thus often rely on indiscriminate web scraping to collect vast data, including personal data, from across the public internet. The guiding assumption is that more data translates into better model accuracy, contextual awareness, and adaptability. Yet this assumption introduces a structural opposition to data minimization, as such models are built not to minimize data input, but to maximize informational capture and generalization capabilities.

Notably, most of the leading LLMs are developed by entities headquartered outside the European Union. This raises additional challenges regarding jurisdiction, enforcement, and the extraterritorial applicability of the GDPR. European users’ personal data may be processed by non-EU actors who do not fully internalize the normative and legal obligations set forth by EU law. Consequently, European Data Protection Authorities (DPAs) face a growing imperative to assert the relevance of EU data protection principles in global technological contexts.

In response to these challenges, the European Data Protection Board (EDPB) has increasingly called for a more expansive interpretation of DPA responsibilities. In its Statement 3/2024, issued in July 2024, the EDPB clarified that DPAs are not merely reactive regulators, but also proactive advisors, coordinators, and ethical arbiters under the forthcoming AI Act[3]. This reconceptualization underscores the need for DPAs to engage not only in enforcement, but in strategic governance and anticipatory oversight of AI technologies.

The present article aims to explore the implications of this evolving regulatory landscape, with a particular focus on the tensions between data minimization and model maximization. It interrogates how DPAs can meaningfully safeguard data protection principles in an era where data volume, rather than data discipline, is increasingly seen as a marker of technological success. Through legal analysis,

[1] *Kuner C., Bygrave L. A., Docksey C.,* The EU General Data Protection Regulation (GDPR): A Commentary, Oxford University Press, 2020.

[2] *Vaezi A.,* Legal Challenges in the Deployment of Large Language Models: A Comparative Analysis under the GDPR and EU AI Act, 2025.

[3] European Data Protection Board, Statement 3/2024 on the role of DPAs under the AI Act, 2024, <https://www.edpb.europa.eu/our-work-tools/our-documents/statements/statement-32024-data-protection-authorities-role-artificial_en> [23.07.2025].

regulatory interpretation, and consideration of emerging enforcement practices, this study contributes to an urgent conversation about the future of responsible AI in the European Union and beyond.

2. Structural Incompatibilities between Data Minimization and AI Model Development

The rapid evolution of large-scale artificial intelligence systems has brought to the forefront a foundational question in contemporary data protection law: Are we witnessing an unavoidable conflict between legal principles and technological practice? At the core of this dilemma lies a fundamental tension between the GDPR's principle of data minimization and the operational architecture of modern machine learning models, particularly large language models (LLMs).

The GDPR, specifically Article 5(1) (c), establishes the principle of data minimization as a cornerstone of lawful data processing[4]. This principle requires that personal data collected must be adequate, relevant, and limited to what is strictly necessary for the specific purposes for which it is processed. In practice, this means that organizations must refrain from collecting or retaining data unless it can be clearly justified in terms of purpose and necessity.

In stark contrast, the logic underlying LLM development is premised on data abundance. The prevailing assumption among AI developers is that the performance and generalizability of these models improve with the volume and diversity of training data. Consequently, LLMs are typically trained on massive datasets encompassing billions of text samples—ranging from academic publications to forum posts, blogs, social media content, and other publicly accessible sources. The aspirational goal is comprehensive linguistic coverage and semantic richness, but this data-centric philosophy directly challenges the necessity and proportionality constraints imposed by the GDPR.

This structural conflict manifests in several critical ways. First, the practice of indiscriminate web scraping often lacks a narrowly defined purpose compatible with data minimization. The mere assumption that all accessible textual data may contribute to model improvement is insufficient under EU data protection law, which requires specific and legitimate processing aims. Furthermore, the scale of collection typically far exceeds what would be considered necessary for the stated objectives of the model, particularly when personal data is involved.

The European Data Protection Board (EDPB), in its Opinion 28/2024, has underscored the importance of rigorous necessity assessments. The Board maintains that personal data embedded within training datasets—even when not directly

[4] *Kuner C., Bygrave L. A., Docksey C.,* The EU General Data Protection Regulation (GDPR): A Commentary, Oxford University Press, 2020.

identifiable—may give rise to re-identification risks[5]. Deep learning models, by virtue of their architecture, can memorize and reproduce training data verbatim or in paraphrased form. This introduces a latent threat that personal information, once included in the training corpus, may later be exposed through model outputs, even in the absence of deliberate intent by the developer.

The EDPB has also drawn attention to the inadequacy of generic anonymization claims. Assertions that personal data has been sufficiently de-identified or pseudonymized must be substantiated by model-specific evaluations and cannot rely on abstract technical assumptions. Inference attacks, model inversion, and membership inference techniques have demonstrated that anonymized data may, under certain conditions, be reverse-engineered or linked back to individuals. These risks necessitate a cautious, case-by-case analysis of the technical safeguards employed in model training.

Moreover, the principle of proportionality is central to the legality of data processing. Controllers must establish a defensible relationship between the quantity of personal data processed and the benefit sought through AI system performance. In the case of LLMs, however, the boundaries of necessity and proportionality are often blurred. Developers frequently fail to articulate why a specific dataset size or composition is required, or why more targeted and privacy-preserving alternatives were not pursued.

The conflict is therefore not merely theoretical but deeply practical. AI developers operate within a paradigm that rewards maximal data ingestion, while data protection frameworks demand restraint, justification, and user-centric safeguards. Bridging this gap will require not only regulatory clarity and enforcement, but also a paradigm shift in how AI innovation is conceptualized.

To move toward compatibility, AI development must increasingly integrate the principles of privacy by design and privacy by default, as mandated under Article 25 of the GDPR. This entails embedding data minimization logic into the architecture of AI systems from their inception. It also implies adopting methodologies that reduce dependency on personal data—such as synthetic data generation, federated learning, or differential privacy mechanisms—thereby aligning technological advancement with legal obligations.

In conclusion, the perceived dichotomy between data minimization and model maximization is emblematic of broader governance challenges in the digital age. While not inherently irreconcilable, these opposing logics require deliberate reconciliation through multidisciplinary collaboration, technical innovation, and regulatory vigilance. Without this effort, the integrity of fundamental rights may be undermined by the unchecked pursuit of technological optimization.

[5] European Data Protection Board, Opinion 28/2024 on Training Data for LLMs, 2024, <https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en> [23.07.2025].

3. Publicly Accessible Data and the Legal Status of Personal Information under the GDPR

A pervasive assumption among AI developers may be that publicly available data is inherently exempt from the full scope of data protection regulation. This misconception underpins many arguments defending the collection and use of large-scale datasets for training AI models, including large language models (LLMs). Developers often argue that because the data used is already accessible on the open web, it does not fall within the regulatory protections of the General Data Protection Regulation (GDPR)[6]. However, such reasoning is legally and ethically flawed.

Under the GDPR, the status of data as "personal" is determined by its identifiability, not its availability. Article 4(1) of the GDPR clearly defines personal data as any information relating to an identified or identifiable natural person, regardless of whether that information was obtained from private or public sources. Thus, the fact that personal data appears on public forums, social media, comment sections, or other open-access domains does not strip it of its protected status.

This legal position has profound implications for the development of AI systems, particularly those trained on data scraped indiscriminately from the internet. The European Data Protection Board (EDPB) has expressed concern about this issue in several recent opinions, most notably in Opinion 28/2024. The EDPB emphasizes that the legality of using such data for training AI models must be assessed on a case-by-case basis. A key point is that the process of data anonymization—often cited by developers as a compliance measure—cannot be presumed effective without rigorous, model-specific validation.

Even if an AI model is not designed to output personal data directly, there remains a significant risk that personal information embedded in the training corpus may be retained within the model's parameters. This latent data may be unintentionally reconstructed in response to user prompts, thereby creating the potential for privacy violations through inference or re-identification. Recent academic studies and real-world incidents have demonstrated that LLMs can inadvertently reproduce specific personal details, such as names, addresses, or fragments of private conversations, raising serious questions about the sufficiency of standard anonymization techniques in the AI context.

In light of these risks, DPAs are increasingly scrutinizing claims of anonymization and demanding transparency about training data composition. To properly assess compliance, regulators must have access to detailed technical documentation, including dataset sources, preprocessing methods, and mitigation strategies. This reinforces the need for regulatory bodies to invest in technical expertise and cross-disciplinary capacity building. Without such capabilities, DPAs cannot perform the

[6] Opinions presented during meeting with Polish DPA. Cf. <uodo.gov.pl> for more information on the meetings with OpenAI or Microsoft.

granular assessments necessary to evaluate whether data minimization, necessity, and proportionality standards have been met.

Furthermore, the principle of privacy by design, articulated in Article 25 of the GDPR, requires that data protection safeguards be embedded into processing activities from the outset. This entails conducting thorough data protection impact assessments (DPIAs) before the commencement of training operations, with a clear articulation of the purpose, scope, and limitations of data collection. Developers must explicitly define what categories of personal data are essential to achieve a model's objectives and demonstrate that less invasive alternatives were considered.

Equally important is the concept of proportionality, which demands a demonstrable relationship between the quantity and sensitivity of personal data processed and the legitimate aims pursued. Massive and indiscriminate scraping of online content—especially without contextual filtering or consent—raises substantial doubts about proportionality and undermines user trust in digital ecosystems.

The idea that "public equals permissible" must be unequivocally rejected. The mere fact that information is accessible online does not confer a license to repurpose it for machine learning without appropriate legal and ethical safeguards. It is a duty of DPAs to challenge this norm and to reinforce the distinction between visibility and validity in data governance.

In conclusion, the lawful use of public data in AI development is far from a settled issue. It calls for robust legal interpretation, rigorous technical scrutiny, and a proactive regulatory posture to ensure that individual rights are not subordinated to the imperatives of technological expansion.

4. Legal Bases for Processing in AI Model Training: Consent, Legitimate Interests, and the Challenge of Transparency

Establishing a valid legal basis for the processing of personal data used in training AI models—particularly large language models (LLMs)—is one of the most complex and disputed issues in the current regulatory landscape. Despite the growing reliance on massive data corpora for developing AI capabilities, many developers have not clearly articulated how such processing complies with the General Data Protection Regulation (GDPR), especially in light of Articles 6, 7, 13, and 14.

While consent is often heralded as the gold standard for lawful processing under the GDPR, its practical application in the AI training context is fraught with difficulty. Article 4(11) defines consent as a freely given, specific, informed, and unambiguous indication of the data subject's wishes. However, this standard is rarely met when consent is sought through general terms of service, opaque privacy policies, or platform-level notices. When personal data is scraped from public websites or user-generated content platforms, there is typically no meaningful opportunity for data subjects to give or withhold consent—let alone understand how their data will be repurposed in AI systems. As such, reliance on consent in these contexts is often legally insufficient and ethically dubious.

In practice, many developers turn to the legal basis of legitimate interest under Article 6(1)(f) as a more flexible alternative[7]. However, the threshold for invoking this basis is stringent. The controller must conduct a comprehensive three-part balancing test: (1) identify a legitimate interest pursued by the data controller or a third party, (2) demonstrate that the processing is necessary for achieving that interest, and (3) prove that the interest is not overridden by the rights and freedoms of the data subject. According to the EDPB's 2024 Guidelines on AI and data processing, this assessment must be both objective and evidence-based, and must be supported by documented safeguards, accountability measures, and mitigation strategies for data subject risk.

A central difficulty in this approach is that necessity and proportionality are rarely self-evident in the context of LLM development. Given the scale and opacity of data scraping, and the speculative nature of benefits derived from diverse training data, it is often unclear whether processing is strictly necessary for the stated purpose or merely convenient for maximizing model performance. The burden of proof lies with the controller to explain why alternative, less invasive methods could not achieve similar results.

Compounding these challenges is the obligation to ensure transparency under Articles 13 and 14 of the GDPR. These provisions require data controllers to inform data subjects about the collection and use of their personal data—whether obtained directly or indirectly. In the context of AI training based on large-scale scraping from multiple platforms, fulfilling this obligation becomes nearly impossible. Developers rarely have access to the identities or contact information of individuals whose data was included in training sets, and retroactive notification is operationally unfeasible.

Nevertheless, the GDPR does not offer an exception to transparency obligations on the basis of scale or technical impracticality. In the absence of effective transparency mechanisms, the lawfulness of the underlying data processing is undermined. This has profound implications for developers seeking to rely on legitimate interest: if affected individuals are not informed, their ability to exercise their rights—such as the right to object under Article 21—is compromised, further weakening the legitimacy of the processing activity.

Moreover, transparency is not only a legal requirement, but also a vital ethical and societal imperative. Trust in AI systems—and in the institutions that govern them—depends on the ability of individuals to understand how their data is being used, and to retain some measure of control over that use. The opacity of many LLMs,

[7] *Sangaraju V. V.,* AI and Data Privacy in Healthcare: Compliance with HIPAA, GDPR, and Emerging Regulations, International Journal of Emerging Trends in Computer Science and Information Technology, 2025, 67–74.

both in terms of their training data and their functioning, exacerbates a broader accountability gap that undermines democratic oversight and public confidence.

In conclusion, the legal bases most commonly invoked for AI training—consent and legitimate interest—are both highly problematic in practice[8]. Developers must go beyond superficial compliance and engage with the spirit of data protection law by embedding transparency, accountability, and necessity into the design and deployment of AI systems. Without such efforts, reliance on these legal grounds may not only fail to meet regulatory scrutiny, but also erode the legitimacy of AI development in the eyes of the public and policymakers alike.

5. Enforcement Trajectories and Strategic Regulatory Responses to AI Data Practices

The enforcement of GDPR provisions in the realm of AI development represents one of the most demanding areas of contemporary data protection. The complexity of AI systems, especially large language models (LLMs), presents regulators with multifaceted challenges, including technical opacity, globalized data flows, and cross-jurisdictional accountability gaps. Despite these barriers, there has been a notable evolution in the posture of European data protection authorities (DPAs), who are increasingly moving from reactive enforcement to coordinated and proactive oversight.

One of the most significant enforcement milestones occurred in 2023, when the Italian Garante per la protezione dei dati personali imposed a temporary ban on ChatGPT. The decision was based on multiple grounds, including the lack of a lawful basis for data processing, insufficient transparency, and the absence of mechanisms to enable data subjects to exercise their rights. This marked the first high-profile intervention by a European DPA against a foundation model, signaling that large-scale AI systems are not immune to GDPR enforcement[9].

Other DPAs have followed suit with both enforcement and guidance. France's CNIL has issued comprehensive recommendations on web scraping, emphasizing that public accessibility does not equate to legal permissibility[10]. The UK's Information Commissioner's Office (ICO) similarly published guidelines articulating the conditions under which AI developers can legally use publicly sourced data for training

[8] *Hoofnagle C. J., van der Sloot B., Zuiderveen Borgesius F.,* The European Union General Data Protection Regulation: What it is and what it means. Information & Communications Technology Law, 28(1), 2019, 65–98.

[9] Garante per la protezione dei dati personali, Decision on OpenAI (ChatGPT), 2023, <https://gdprhub.eu/index.php?title=Garante_per_la_protezione_dei_dati_personali_(Italy)_-_9870832> [23.07.2025].

[10] CNIL, First Recommendations on the Development of AI Systems, Commission Nationale de l'Informatique et des Libertés, 2024, <https://www.cnil.fr/en/ai-cnil-publishes-its-first-recommendations-development-artificial-intelligence-systems> [23.07.2025].

purposes[11]. Meanwhile, Greece's Hellenic Data Protection Authority imposed a €20 million fine on Clearview AI, establishing a precedent for holding biometric data scrapers accountable under GDPR provisions[12]. These actions illustrate the increasing willingness of DPAs to challenge powerful AI actors.

A key development at the European level is the emergence of joint enforcement mechanisms. The European Data Protection Board (EDPB) has launched dedicated task forces—most notably those concerning ChatGPT and DeepSeek[13]—that enable coordinated investigations and harmonized interpretations across Member States. These efforts are likely to be institutionalized under the AI Act, which introduces an EU-wide AI governance framework featuring the European Artificial Intelligence Office and enhanced cross-border supervisory structures. These new arrangements echo the GDPR's "One-Stop-Shop" model but with a stronger emphasis on systemic risk assessments and sector-specific oversight.

In parallel, DPAs are beginning to articulate forward-looking regulatory strategies. These include issuing proactive guidelines, demanding algorithmic impact assessments, and exploring technical audit procedures for model explainability and training data lineage. The trajectory is clear: enforcement is no longer confined to penalizing past violations but now encompasses ex ante regulation designed to prevent systemic harms before they materialize.

6. Ethical Oversight in AI Development: Expanding the Role of Data Protection Authorities

While legal frameworks such as the GDPR and the upcoming AI Act provide formal criteria for compliance, they are not always equipped to fully address the ethical dimensions of AI development. The use of personal data to train generative AI models invokes broader societal concerns related to human dignity, individual autonomy, and cultural representation. Practices that are technically lawful under a narrow interpretation of the law may still provoke ethical unease, public backlash, or social harm.

One salient example is the use of expressive personal data—such as voice recordings, biometric images, or creative content—to generate synthetic media. While developers may argue that such uses fall within lawful grounds if the data was publicly accessible, this overlooks the deeper issue of consent, artistic ownership, and the right

[11] ICO, The Lawful Basis for Web Scraping to Train Generative AI Models, Information Commissioner's Office (UK), 2024, <https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/response-to-the-consultation-series-on-generative-ai/the-lawful-basis-for-web-scraping-to-train-generative-ai-models/> [23.07.2025].

[12] Info on: <https://www.edpb.europa.eu/news/national-news/2022/hellenic-dpa-fines-clearview-ai-20-million-euros_en> [23.07.2025].

[13] *Deng Z., Ma W., Han Q. L., Zhou W., Zhu X.,* Exploring DeepSeek: A Survey on Advances, Applications, Challenges and Future Directions, IEEE/CAA Journal of Automatica Sinica, 12(5), 2025, 872–893.

to one's likeness. In creative and journalistic sectors, unauthorized use of archival material to train AI-generated voices or avatars has been widely condemned as exploitative. Such practices erode professional integrity and diminish the ability of individuals to control how their identities are digitally reproduced.

These concerns extend beyond individual harms to systemic risks, including deepfakes, disinformation, political manipulation, and cultural homogenization. Synthetic content generated by LLMs or multimodal AI systems can be used to imitate real individuals, skew public discourse, or undermine democratic processes. The ethical implications of such uses are profound and not always foreseeable at the point of data collection or model training.

Recognizing these risks, the EDPB has called upon DPAs to assume a broader ethical mandate. Beyond ensuring formal compliance, DPAs are increasingly expected to identify power asymmetries between individuals and developers, evaluate the societal impact of large-scale data exploitation, and promote responsible innovation. Ethical assessment should therefore be treated as a complementary dimension of data protection governance, integrated into risk-based regulation and transparency obligations.

To fulfill this role effectively, DPAs must develop interdisciplinary capacities—encompassing legal, technical, sociological, and philosophical expertise—and cultivate dialogue with affected communities, civil society organizations, and academic researchers. Ethical oversight should be embedded in algorithmic design through mechanisms such as fairness audits, participatory model evaluations, and public interest impact statements.

In sum, ethical stewardship is emerging as a critical extension of the regulatory function. It repositions DPAs not only as guardians of legality but as arbiters of justice in a rapidly evolving technological landscape. Only by aligning legal compliance with ethical legitimacy can the governance of AI systems uphold public trust and fundamental rights in the digital era.

7. Future Directions: Integrating the AI Act and Addressing Institutional Challenges

The – already mentioned - European Union Artificial Intelligence Act (AI Act) represents a landmark regulatory initiative that introduces a comprehensive, risk-based framework for governing AI systems. Building upon the foundations established by the GDPR, the AI Act imposes heightened obligations for so-called high-risk AI applications. These include, among others, mandatory conformity assessments, structured technical documentation, robust data governance frameworks, and post-market monitoring requirements.

One of the most significant contributions of the AI Act is the operationalization of data quality and data minimization principles within the AI lifecycle. Article 10 (3) of the AI Act explicitly mandates that datasets used for training, validation, and testing must be "relevant, representative, free of errors, and as complete as possible," while also emphasizing that they should be minimized in scope to avoid unnecessary exposure of personal data. In this sense, the Act strengthens the normative trajectory initiated by the GDPR by embedding data protection standards directly into AI system design and evaluation.

For Data Protection Authorities (DPAs), the implementation of the AI Act signifies both an expansion of responsibilities and a transformation of institutional identity. DPAs will no longer act solely as national enforcers of data privacy but must evolve into key nodes in a pan-European network of AI governance. This includes collaboration with the newly established European Artificial Intelligence Office, participation in cross-border investigations, contribution to harmonized guidance, and oversight of high-risk AI systems deployed across multiple sectors, including health, education, employment, and law enforcement.

However, this transition is fraught with challenges. Many DPAs currently face substantial resource limitations, including understaffing, limited technical infrastructure, and insufficient in-house expertise in machine learning, algorithmic auditing, and systems engineering. The new responsibilities outlined in the AI Act—such as the capacity to evaluate training data lineage, assess algorithmic impact, and ensure conformity with design-level transparency—will require significant investment in organizational capacity, skills development, and institutional coordination.

Another source of complexity arises from the doctrinal and operational intersections between the GDPR and the AI Act. Developers must navigate overlapping, and at times potentially conflicting, obligations concerning data minimization, lawful basis for processing, fairness, accountability, and data subject rights. These overlaps will necessitate interpretative guidance from the EDPB and the European AI Office to ensure coherent and non-redundant enforcement. The development of joint compliance frameworks and model templates may help to bridge regulatory gaps and promote legal certainty for developers operating across multiple EU jurisdictions.

Finally, the global nature of AI development poses challenges to the enforcement reach of European regulations. Many foundational models are developed outside the EU, and their integration into local products or services often obscures jurisdictional boundaries. The success of the AI Act will depend on the ability of European regulators to assert extraterritorial influence through cooperation mechanisms, adequacy frameworks, and public procurement incentives that favor compliant systems.

In sum, the AI Act offers an unprecedented opportunity to align technological innovation with democratic values and fundamental rights. Yet its implementation will require robust institutions, cross-sectoral cooperation, and sustained political commitment to make responsible AI not just a regulatory aspiration, but an operational reality.

8. Conclusion: From Legal Compliance to Ethical and Strategic Stewardship

The principle of data minimization, once seen as a technical constraint or bureaucratic formality, has emerged as a normative bulwark against surveillance capitalism, algorithmic exploitation, and asymmetries of power in the digital era. In an age increasingly defined by model maximization and data commodification, it serves as both a legal requirement and a moral imperative.

Yet the application of this principle must evolve in response to the unique complexities posed by contemporary AI systems. Large language models and other foundation models challenge conventional legal categories and procedural safeguards, calling for a more dynamic and holistic approach to regulatory enforcement. As such, Data Protection Authorities must reconceptualize their mandate—not only enforcing compliance, but fostering systemic accountability, ethical reflection, and public trust.

This expanded role entails resisting unjustified or disproportionate data practices, promoting transparent and explainable AI, and safeguarding the rights and freedoms of individuals whose data underpins digital innovation. It also requires building institutional capacity to conduct risk-based audits, engage with civil society, and contribute to the ethical governance of AI technologies.

Ultimately, the responsible development and deployment of AI cannot be reduced to a checklist of legal obligations. It is a collective societal commitment to embedding human dignity, fairness, and justice at the core of technological progress. In this endeavor, DPAs are not just regulators—they are stewards of the digital public interest.

Bibliography:

1. CNIL, First Recommendations on the Development of AI Systems, Commission Nationale de l'Informatique et des Libertés, 2024, <https://www.cnil.fr/en/ai-cnil-publishes-its-first-recommendations-development-artificial-intelligence-systems> [23.07.2025].
2. *Deng Z., Ma W., Han Q. L., Zhou W., Zhu X.,* Exploring DeepSeek: A Survey on Advances, Applications, Challenges and Future Directions, IEEE/CAA Journal of Automatica Sinica, 12(5), 2025, 872–893.
3. European Data Protection Board, Statement 3/2024 on the role of DPAs under the AI Act, 2024, <https://www.edpb.europa.eu/our-work-tools/our-documents/statements/statement-32024-data-protection-authorities-role-artificial_en> [23.07.2025].
4. European Data Protection Board, Opinion 28/2024 on Training Data for LLMs, 2024 <https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en> [23.07.2025].
5. Garante per la protezione dei dati personali, Decision on OpenAI (ChatGPT), 2023, <https://gdprhub.eu/index.php?title=Garante_per_la_protezione_dei_dati_personali_(Italy)_-_9870832> [23.07.2025].
6. *Hoofnagle C. J., van der Sloot B., Zuiderveen Borgesius F.,* The European Union General Data Protection Regulation: What it is and what it means. Information & Communications Technology Law, 28(1), 2019, 65–98.
7. ICO, The Lawful Basis for Web Scraping to Train Generative AI Models, Information Commissioner's Office (UK), 2024, <https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/response-to-the-consultation-series-on-generative-ai/the-lawful-basis-for-web-scraping-to-train-generative-ai-models/> [23.07.2025].
8. *Kuner C., Bygrave L. A., Docksey C.,* The EU General Data Protection Regulation (GDPR): A Commentary, Oxford University Press, 2020.
9. *Sangaraju V. V.,* AI and Data Privacy in Healthcare: Compliance with HIPAA, GDPR, and Emerging Regulations, International Journal of Emerging Trends in Computer Science and Information Technology, 2025, 67–74.
10. *Vaezi A.,* Legal Challenges in the Deployment of Large Language Models: A Comparative Analysis under the GDPR and EU AI Act, 2025.