



## AI chatbots meet European privacy enforcement

Rory Macmillan<sup>1</sup> and Juan Abad

Large Language Models (“LLMs”) have challenged the legal world since their public roll-out. DeepSeek’s global debut in January 2025 was no exception. For the first time, the Western world experienced the deployment of an LLM built and trained in the People’s Republic of China.<sup>2</sup> Against a background of rising geopolitical tensions between the East and the West, fears relating to “artificial general intelligence”<sup>3</sup>, and growing privacy concerns, European data protection authorities stepped up their engagement.

European data regulators had already scrutinized LLMs of developers from Silicon Valley, including Italy’s suspension of ChatGPT in 2023.<sup>4</sup> These precedents revealed how Europe’s enforcers were testing the limits of the General Data Protection Regulation (“GDPR”) and the emerging Artificial Intelligence (“AI”) Act to keep pace with LLMs.

This article examines the rise of enforcement actions taken by European data authorities and the privacy dilemmas posed by LLMs under the GDPR. It considers how these efforts intersect with the EU’s AI Act, the geopolitical implications of cross-border data flows, and the remedies and coordination challenges that shape enforcement today. Finally, it offers observations about how Europe’s evolving regulatory framework may influence the global trajectory of AI governance.

### 1. *The Rise of LLMs*

The launch of ChatGPT in late 2022 triggered a new wave of scrutiny over how AI systems handle personal data.<sup>5</sup> In March 2023, Italy’s data protection authority, the *Garante*, suspended ChatGPT’s operations in Italy, citing concerns that OpenAI lacked a clear legal basis to process

---

<sup>1</sup> Rory Macmillan is a founding partner of Macmillan Keck, Attorneys & Solicitors and serves as Vice-Chair of the Privacy and Information Security Committee of the American Bar Association. The author thanks Gabriel Obando-Chacón and Kemal Naqvi for their research assistance.

<sup>2</sup> Laura Caroli, *DeepSeek: A Problem or an Opportunity for Europe?*, *Ctr. for Strategic & Int’l Studies* (Feb. 14, 2025), <https://www.csis.org/analysis/deepseek-problem-or-opportunity-europe>.

<sup>3</sup> *Pause Giant AI Experiments: An Open Letter*, Future of Life Institute (Mar. 22, 2023), <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.

<sup>4</sup> Nick Cote & Sheera Frenkel, *ChatGPT, the AI Tool, Was Temporarily Banned in Italy*, *N.Y. Times* (Mar. 31, 2023), <https://www.nytimes.com/2023/03/31/technology/chatgpt-italy-ban.html>.

<sup>5</sup> Bernard Marr, *A Short History of ChatGPT: How We Got to Where We Are Today*, *Forbes* (May 19, 2023), <https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/>.



personal data for model training.<sup>6</sup> The *Garante* also ordered the company to clarify how it collected, stored and corrected information about individuals, to adopt age verification tools, and to give users an accessible way to exercise their rights under the GDPR.<sup>7</sup> When OpenAI later failed to fully comply, Italy imposed a €15 million (approximately USD 16 million) fine in December 2024.<sup>8</sup>

The focus then shifted to Ireland, which had become OpenAI's European hub under the European Union's "one-stop-shop" system for investigations.<sup>9</sup> Ireland's Data Protection Commission ("DPC") launched inquiries into ChatGPT and other major players. In mid-2024, the Irish DPC initiated proceedings against X (formerly Twitter) for using European users' data to train its AI chatbot, Grok, without consent.<sup>10</sup> After negotiations, X agreed to stop the practice, but still faced risk that failure to delete previously collected data could still trigger penalties.<sup>11</sup> At the same time, the Irish DPC began examining whether Google had conducted a Data Protection Impact Assessment before developing its large-scale language model, PaLM2.<sup>12</sup> The case went beyond reacting to public complaints in probing the internal design and accountability processes behind AI systems.

The arrival of DeepSeek in early 2025 reignited these concerns on a larger scale. Developed and trained in China, DeepSeek's chatbot quickly gained traction across Europe before being blocked in Italy for failing to address privacy risks.<sup>13</sup> Within weeks, regulators in France, Ireland, Belgium, Greece and the Netherlands<sup>14</sup> opened parallel investigations.<sup>15</sup> Their questions focused on whether the company's transfer of European user data to servers in China could ever comply

---

<sup>6</sup> *Provvedimento del 30 marzo 2023*, Garante per la protezione dei dati personali (Mar. 30, 2023), <https://www.gpdp.it/web/guest/home/docweb/-/docweb-display/docweb/9870832>.

<sup>7</sup> *Id.*

<sup>8</sup> *Provvedimento del 3 luglio 2024*, Garante per la protezione dei dati personali (July 3, 2024), <https://gpdp.it/web/guest/home/docweb/-/docweb-display/docweb/10085432>.

<sup>9</sup> Natasha Lomas, *OpenAI Moves to Shrink Regulatory Risk in EU Around Data Privacy*, TechCrunch (Jan. 2, 2024), <https://techcrunch.com/2024/01/02/openai-dublin-data-controller/>.

<sup>10</sup> *X Permanently Stops Grok AI from Using EU Citizens' Tweets After Court Action by Irish Data Watchdog*, Indep. (Sep. 4, 2024), <https://www.independent.ie/business/technology/x-permanently-stops-grok-ai-from-using-eu-citizens-tweets-after-court-action-by-irish-data-watchdog/a168142842.html>.

<sup>11</sup> *Id.*

<sup>12</sup> *Top EU Privacy Regulator Opens Probe into Google's AI Compliance*, Reuters (Sept. 11, 2024), <https://www.reuters.com/technology/artificial-intelligence/top-eu-privacy-regulator-opens-probe-into-googles-ai-compliance-2024-09-11/>.

<sup>13</sup> *Italy's Privacy Watchdog Blocks Chinese AI App DeepSeek*, Reuters (Jan. 30, 2025), <https://www.reuters.com/technology/artificial-intelligence/italys-privacy-watchdog-blocks-chinese-ai-app-deepseek-2025-01-30/>.

<sup>14</sup> The Dutch Data Protection Authority not only launched a formal investigation but also issued a public advisory warning users against sharing sensitive data with the chatbot, reflecting a more preventive enforcement posture.

<sup>15</sup> *DeepSeek Faces Expulsion from App Stores in Germany*, Reuters (June 27, 2025), <https://www.reuters.com/sustainability/boards-policy-regulation/deepseek-faces-expulsion-app-stores-germany-2025-06-27/>.



with the GDPR, given that China is not recognized as offering adequate data protection. The European Data Protection Board (“EDPB”) expanded its AI taskforce to coordinate these national efforts and warned that DeepSeek and similar AI systems would face heightened scrutiny going forward.<sup>16</sup>

These developments revealed a shift from seeking to understand how the law applies to generative AI to a focus on enforcement. Regulators are no longer debating whether AI models fall under the GDPR. They are now determining whether their data collection, training and deployment practices conform with the regulations in effect and applying enforcement where they do not.

## 2. *The core privacy dilemma*

At the heart of Europe’s regulatory debate lies what appears to be a simple question: how can AI systems learn from human data without violating data protection laws? When EDPB published its long-awaited opinion in December 2024, it exposed the tension between the data-hungry nature of LLMs and the boundaries of the GDPR. The opinion, requested by Ireland’s DPC, distilled years of fragmented discussions into a clear message: most AI training data is not anonymous, and claiming otherwise does not make it so.

The EDPB emphasized that anonymity under the GDPR sets a high bar. Developers frequently argue that scraped or aggregated datasets are “de-identified”, yet generative models can reproduce fragments of those same datasets word for word (a phenomenon referred to as “regurgitation”).<sup>17</sup>

In the Board’s view, data is only truly anonymized when the chance of re-identification is insignificant. In practice, few AI systems meet that test. Even partial or probabilistic reconstruction of personal data can bring an entire model back within the GDPR’s scope. For developers, this means that most training datasets remain subject to the GDPR’s full compliance obligations.<sup>18</sup>

Beyond anonymization, the EDPB turned its attention to the lawful basis of processing. AI developers often rely on “legitimate interest” when consent is impractical. Under the GDPR’s Article 6(1)(f), controllers must show three elements: a clear and lawful purpose, the necessity of the processing for that purpose, and a fair balance between the organization’s goals and the rights of individuals. Examples of legitimate purposes include detecting fraud, reducing algorithmic bias, and improving system security. Yet the balancing test is delicate. Processing sensitive information

---

<sup>16</sup> *EDPB Adopts Statement on Age Assurance, Creates Task Force on AI Enforcement, and Gives Opinion on GDPR Harmonisation Proposal*, European Data Protection Bd. (Jan. 23, 2025), [https://www.edpb.europa.eu/news/news/2025/edpb-adopts-statement-age-assurance-creates-task-force-ai-enforcement-and-gives\\_en](https://www.edpb.europa.eu/news/news/2025/edpb-adopts-statement-age-assurance-creates-task-force-ai-enforcement-and-gives_en).

<sup>17</sup> *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*, European Data Protection Bd. (Dec. 17, 2024), [https://www.edpb.europa.eu/system/files/2024-12/edpb\\_opinion\\_202428\\_ai-models\\_en.pdf](https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf).

<sup>18</sup> *Id.*



such as financial or health data will rarely pass because potential harm — from identity theft to reputational damage — can outweigh any commercial or technical benefit.<sup>19</sup>

The Board also highlighted “reasonable expectations”.<sup>20</sup> People might tolerate their public posts being used to train an AI summarizer, but not private messages or internal documents collected without notice.

For AI companies, this generally leaves consent and legitimate interest as the only viable legal routes, each with limits. Consent offers clarity but is almost impossible to scale across billions of data points. Legitimate interest is scalable but fragile under regulatory scrutiny. Italy’s fine against OpenAI underscored that a poorly documented balancing test can itself be a violation. The EDPB has urged firms to record their reasoning in detail, adopt pseudonymization or differential-privacy techniques, and design visible opt-out mechanisms for users.<sup>21</sup>

A further question addressed in the opinion concerned downstream use. What happens if a model trained on unlawfully obtained data is later deployed by another organization? The Board concluded that deployment is not automatically unlawful so long as the deployer does not continue the underlying illegal processing. Hospitals using an AI diagnostic tool, for example, would not be liable unless they knew the model had been trained on non-compliant datasets. Still, they must perform due diligence and seek contractual assurances from suppliers. This might be an early sign that AI governance will increasingly extend across the entire supply chain.<sup>22</sup>

Together, these principles define Europe’s emerging privacy dilemma. The continent is not attempting to halt innovation but to ensure that innovation is explainable, traceable and lawful. By clarifying that most AI data remains personal data, and by insisting on accountability across development and deployment, the EDPB drew a new line in the sand. For the global AI industry, Europe’s message was that training smarter models must not come at the expense of privacy.

### *3. The AI Act and GDPR: overlaps and tensions*

While the GDPR remains the backbone of Europe’s data protection regime, it is no longer the only law shaping how AI must be designed and deployed. With the EU’s AI Act adopted in August 2024 and coming into force in stages, regulators and developers now find themselves navigating two parallel systems of accountability. Both aim to align technology to European legal values, yet they differ in focus and method.

---

<sup>19</sup> *Id.*

<sup>20</sup> *Id.*

<sup>21</sup> *Id.*

<sup>22</sup> *Id.*



The GDPR governs how personal data are collected and processed, establishing lawful bases, individual rights, and organizational duties to ensure data protection by design.<sup>23</sup> The AI Act, by contrast, regulates the development and deployment of AI systems themselves, setting standards for design, safety, and risk management across varying levels of potential harm.<sup>24</sup>

The AI Act introduces a risk-based framework that classifies systems according to the level of harm they can cause. It prohibits practices considered incompatible with European values, such as social scoring and untargeted facial recognition. It then imposes strict obligations on “high-risk” systems — those used in employment, healthcare, education, or law enforcement — requiring risk assessments, human oversight and security testing.<sup>25</sup>

In parallel, the Act introduces a separate category for general-purpose AI models, such as ChatGPT or DeepSeek, which can be deployed across multiple contexts. These models are not automatically considered low risk; rather, they are subject to a dedicated set of transparency and governance duties, including requirements to document training data, ensure cybersecurity, and disclose system capabilities and limitations.<sup>26</sup>

At first glance, the AI Act and the GDPR seem complementary. One focuses on ethics and product safety, the other on privacy and lawful data use. Yet as enforcement begins, points of friction are already emerging. The most significant lies in the handling of sensitive data. The AI Act allows limited processing of such data when it is “strictly necessary” to monitor or correct bias in high-risk systems.<sup>27</sup> Under the GDPR, however, processing health, biometric, or ethnic data requires explicit consent or a public-interest exception.<sup>28</sup> An AI company attempting to improve fairness in recruitment by collecting demographic data might therefore comply with one law while violating another. Regulators have not yet resolved this contradiction, leaving companies to rely on cautious case-by-case analysis.

Another overlap arises around accountability. Both laws require some form of impact assessment, but with different emphasis. The GDPR mandates a Data Protection Impact Assessment when data processing may pose a high risk to individual rights.<sup>29</sup> The AI Act demands a broader conformity assessment, covering technical safety, robustness, and governance processes.<sup>30</sup> Together, they form a double layer of compliance that can easily confuse even well-

---

<sup>23</sup> Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 Apr. 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (GDPR).

<sup>24</sup> Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on Artificial Intelligence (AI Act), 2024 O.J. (L 1689) 1.

<sup>25</sup> AI Act, arts. 5(1)(c)–(d), 6–9, 10, 14–15, Annex III.

<sup>26</sup> *Id.*, arts. 52–56.

<sup>27</sup> *Id.*, art. 10(5).

<sup>28</sup> GDPR, art. 9(1), 9(2)(a), 9(2)(g).

<sup>29</sup> *Id.*, art. 35.

<sup>30</sup> AI Act, arts. 43–49 & Annex VII.



intentioned developers. For multinational firms, integrating these assessments into a single compliance framework has become a pressing challenge.

Still, the convergence between both regimes offers an unexpected advantage. By forcing companies to document and explain how their systems process data, the GDPR acts as a “pre-AI Act” mechanism. Firms that already apply privacy-by-design principles often find themselves better prepared for the AI Act’s transparency and traceability standards. In this sense, the GDPR continues to serve as Europe’s de facto blueprint for responsible AI development, while the AI Act extends those values into the technical fabric of machine learning.

Where the GDPR asks *why* data is used, the AI Act asks *how* the system behaves once the data is used. Together, they are redefining the boundaries between innovation, accountability, and risk. Yet the true test will come in enforcement, as regulators determine whether these two frameworks can coexist without paralyzing the very innovation they aim to guide.

#### 4. *Data without borders: the geopolitics of AI enforcement*

The discussion around AI in Europe has become intertwined with geopolitics. Every conversation about data protection now carries a deeper question: who controls the data that fuels AI models, and under whose laws does that data fall? When DeepSeek entered the European market, these concerns came to the forefront. Its privacy policy expressly stated that user data would be stored on servers in China, immediately raising doubts about whether such transfers could ever meet the GDPR’s requirements.<sup>31</sup>

Under European law, personal data may only be transferred outside the European Economic Area if the destination country ensures an adequate level of protection. China is not recognized as providing that standard. Without an adequacy decision, companies must rely on other legal mechanisms such as contractual safeguards or binding corporate rules. These frameworks, while theoretically available, are difficult to monitor and even harder to enforce when the data recipient operates under a fundamentally different legal and political system. For regulators, this means that an AI service transferring data to China risks violating European privacy rules simply by design. The issue is not only legal but strategic, touching on data sovereignty and the broader security implications of cross-border data flows.

This growing sensitivity is part of a wider international context. In 2024, the United States issued an executive order restricting the export of sensitive personal data to adversarial states.<sup>32</sup>

---

<sup>31</sup> Matt Burgess, *DeepSeek’s Popular AI App Is Explicitly Sending US Data to China*, *Wired* (Jan. 27, 2025), <https://www.wired.com/story/deepseek-ai-china-privacy-data/>.

<sup>32</sup> *Exec. Order No. 14117, Preventing Access to Americans’ Bulk Sensitive Personal Data and United States Government-Related Data by Countries of Concern*, 89 Fed. Reg. 12035 (Feb. 28, 2024),



Europe is pursuing the same concern from a different angle, framing it within the language of privacy and rights rather than national security. Both approaches lead to stricter scrutiny over how and where citizens' data travels, and stronger expectations of transparency from those who process it.

The contrast between data transfers to China and those to the United States highlights how geopolitics shapes privacy enforcement. Since 2023, the EU–U.S. Data Privacy Framework has re-established a legal channel for transatlantic data transfers under defined safeguards.<sup>33</sup> Companies like OpenAI and Google can rely on this mechanism to store and process European data on U.S. servers while remaining within the GDPR's boundaries. No equivalent arrangement exists for China, which is why DeepSeek faced swift regulatory action while American-developed models continue to operate under close but lawful oversight.

Beyond individual cases, Europe's approach to enforcement also reflects its role in global digital governance. Some critics have portrayed the GDPR enforcement as a way of protecting European markets from U.S. technology dominance. Yet this interpretation overlooks the underlying intent. Europe's insistence on rights-based regulation stems from a belief that data protection is not merely a compliance matter but a fundamental expression of democratic values. By setting high standards for accountability, the EU seeks to protect its citizens and influence how technology is built.

Extraterritoriality remains central to this endeavor. The GDPR applies not only to companies established within the EU but also to those outside offering services to, or processing the data of, individuals in Europe. This enhanced scope projects European values beyond the continent's borders, requiring even non-European developers to adapt to its standards. Where data has become both an economic resource and a geopolitical instrument, compliance with European privacy norms has significant implications for credibility, reputation and trust.

##### *5. Building accountability: the next stage of AI enforcement*

European data regulators are no longer content with cautionary guidance. They are acting. While fines remain a familiar deterrent, enforcement has evolved toward measures that reshape how AI systems function. Authorities now suspend processing, restrict model training, and order the deletion of unlawfully sourced datasets. In severe cases, they may require that an entire model

---

<https://www.federalregister.gov/documents/2024/03/01/2024-04573/preventing-access-to-americans-bulk-sensitive-personal-data-and-united-states-government-related>.

<sup>33</sup> Joe Duball, *European General Court Dismisses Latombe Challenge, Upholds EU-US Data Privacy Framework*, *Int'l Ass'n of Privacy Professionals* (Sept. 3, 2025), <https://iapp.org/news/a/european-general-court-dismisses-latombe-challenge-upholds-eu-us-data-privacy-framework/>.



be withdrawn if its foundation cannot be brought into compliance. These tools give the GDPR real operational bite.

Recent decisions illustrate how these powers are being used. Italy's *Garante* imposed a hefty penalty on OpenAI and obliged it to launch a public awareness campaign explaining how user data is handled. The Irish DPC opened proceedings against X to block its chatbot training on European data, while Dutch and French regulators have targeted smaller AI developers for scraping personal information without legal basis.<sup>34</sup> The message is consistent: accountability extends to the design and maintenance of AI systems, not merely to their outputs.

In early 2025, Italy's *Garante* also fined the developer of the Replika chatbot €5 million (approximately USD 5.4 million) for multiple GDPR violations, citing the processing of sensitive data without adequate safeguards.<sup>35</sup> The case broadened the scope of enforcement beyond large foundation models to consumer-facing AI tools, confirming that even smaller operators remain within the reach of European regulators.

To avoid inconsistent outcomes, national authorities are increasingly coordinating their work through the EDPB. The Board's expanded AI taskforce, launched in February 2025, enables joint investigations, common methodologies, and shared expertise across member states. This coordination has already brought coherence to parallel actions concerning DeepSeek, Google's PaLM2, and OpenAI, reducing the risk of fragmented enforcement and strengthening the credibility of Europe's digital oversight.

Still, differences remain in pace and tone. Some regulators, notably in Italy and France, favor rapid intervention. Others, such as Ireland's DPC, proceed with caution. Yet the broader direction is that enforcement is becoming collective, strategic, and increasingly systemic. Investigations now extend beyond AI developers to include data brokers, hosting providers, and third-party suppliers. Compliance depends on the integrity of entire data supply chains.

The next phase may blend privacy enforcement with other areas of European law. Product-liability and consumer-protection reforms under discussion in Brussels could soon expose AI developers to damages claims for harm caused by biased or unsafe systems. As boundaries between data protection, safety and consumer rights blur, companies will be judged not only on their respect for privacy but on whether their technology behaves responsibly in practice.

---

<sup>34</sup> *Data Scraping: French Supervisory Authority Fined KASPR €240,000*, European Data Protection Board (Jan. 23, 2025), [https://www.edpb.europa.eu/news/news/2025/data-scraping-french-supervisory-authority-fined-kaspr-eu240-000\\_en](https://www.edpb.europa.eu/news/news/2025/data-scraping-french-supervisory-authority-fined-kaspr-eu240-000_en).

<sup>35</sup> *AI: the Italian Supervisory Authority Fines Company Behind Chatbot "Replika"*, European Data Protection Board (May 21, 2025), [https://www.edpb.europa.eu/news/national-news/2025/ai-italian-supervisory-authority-fines-company-behind-chatbot-replika\\_en](https://www.edpb.europa.eu/news/national-news/2025/ai-italian-supervisory-authority-fines-company-behind-chatbot-replika_en).



For the AI industry, this signals a shift from defensive compliance towards continuous governance. Real-time dataset audits, privacy-preserving architectures, and transparent accountability chains are increasingly important for market trust. Europe's regulators are influencing a global conversation about what responsible AI looks like.

### *Conclusion*

As AI systems expand in scale and influence, their legitimacy will depend on the principles that guide their creation and use. Europe's approach, grounded in privacy, transparency, and accountability, seeks to demonstrate that technological progress can coexist with the protection of individual rights. Whether it achieves an enduring balance amid global competition will affect the future of AI regulation and the trust underpinning technological development.