

The Open-Weight AI Compliance:

Why “Free” Models
Carry Hidden
Regulatory Costs



ARTICLE

AUTHOR: **TANYA CHIB**

PRIVACY RULES

Organizations downloading open-source AI models for high-risk applications may face the worst of both worlds: full regulatory liability with none of the provider support that proprietary systems offer.

The Deceptive Simplicity

A healthcare startup downloads Meta's LLaMA 3 to screen patient referrals. A fintech company deploys Mistral 7B for loan decisions. A municipal housing authority uses Falcon 40B to rank subsidy applicants. Each scenario has three things in common: the technology is freely available, deployment takes hours, and regulatory liability is profound.

These aren't hypothetical edge cases. These cases are likely to be classified under Annex III high-risk AI applications under Regulation (EU) 2024/1689 (the "**AI Act**")¹, a designation that triggers comprehensive compliance obligations including technical documentation², human oversight³, quality management systems⁴, and fundamental rights impact assessments⁵. The sticky situation being that while model providers enjoy broad regulatory exemptions, deployers of open-weight models in high-risk contexts assume the full compliance burden alone.

The Three-Tier Liability Structure

The AI Act's Article 2(5) creates an intentional asymmetry. Open-source AI components are generally exempt from provider obligations, unless "*placed on the market or put into service as high-risk AI systems.*"⁶ This produces three regulatory tiers:

1 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024.

2 AI Act, Article 11 (Technical documentation).

3 AI Act, Article 14 (Human oversight).

4 AI Act, Article 17 (Quality management system).

5 AI Act, Article 27 (Fundamental rights impact assessment for certain deployers of high-risk AI systems).

6 AI Act, Article 2(5). See also Article 3(11) defining "putting into service" as supply for first use directly to the deployer.

Tier 1: Model providers (Meta, Mistral, Falcon) releasing weights with no deployment, face minimal obligations.

Tier 2: Organizations using open-weight models for low-risk applications (translation, content generation) maintain exemption benefits.

Tier 3: Organizations deploying the same models in high-risk contexts (healthcare, credit, public benefits) lose all exemptions and face obligations equivalent to developing proprietary AI systems from scratch.

As Recital 60 explains, risk emerges from deployment context, not model existence⁷. The same language model that poses minimal risk for translation becomes high-risk when screening medical referrals. The AI Act regulates where risk crystallizes, i.e. at the point of high-risk deployment.

The Information Vacuum

The compliance challenge extends beyond formal obligations to a structural information deficit. Recent empirical analysis of nearly 460,000 AI model cards from HuggingFace reveals a troubling gap, only 14% contain any risk information⁸, and even those systematically underreport deployment-relevant risks. The data speaks volumes:

- **Privacy and security risks:** 3% of developer reports vs. 8% of real-world incidents.⁹
- **Malicious use and misinformation:** 14% of reports vs. 35% of documented harms.¹⁰
- **Technical performance issues:** 37% of reports vs. 24% of incidents.¹¹

Developers appropriately emphasize technical capabilities but largely overlook the privacy, transparency, and human-interaction risks that GDPR and funda-

7 Recital 60: "Free and open-source AI components should not fall within the scope of this Regulation unless they are placed on the market or put into service as high-risk AI systems or as an AI system that falls under the prohibition of prohibited practices."

8 Rao, P. S. B., Scepanovic, S., Jayagopi, D. B., Cherubini, M., & Quercia, D. (2025). The AI Model Risk Catalog: What Developers and Researchers Miss About Real-World AI Harms. Proceedings of the AAAI Conference on Artificial Intelligence, p. 2 (analyzing 461,181 model cards from HuggingFace, of which 64,116 (14%) contained risk-related sections).

9 Id. at p. 6, Figure 3 (model cards show 2.9% privacy & security risks vs. AI Incident Database showing 8.2%).

10 Id. (model cards show 4.0% + 10.2% = 14.2% for malicious actors & misinformation vs. incident database showing 22.4% + 12.9% = 35.3%).

11 Id. (model cards show 37.3% AI system safety, failures & limitations vs. incident database showing 23.9%).

mental rights assessments demand¹². For high-risk deployers, this means upstream documentation systematically underrepresents the compliance risks they must manage.

Interview research confirms the practical impact. One practitioner stated bluntly, *“With open-source models, we have no real insight into how they were trained, which makes it extremely difficult to ensure compliance in high-risk contexts.”*¹³

The GDPR Compliance Challenge

High-risk deployers face two particularly acute GDPR challenges with open-weight models:

Article 17 (Right to Erasure): When an EU resident requests data deletion from a model trained on millions of web-scraped documents, deployers face an impossible choice. Model retraining with specific data removed is technically complex and economically prohibitive. Post-processing filters cannot guarantee complete removal from learned representations. The Spanish Data Protection Authority’s guidance on AI risk assessment suggests organizations must demonstrate capacity to implement data subject rights¹⁴, for high-risk applications, inability to comply may constitute fundamental compliance failure.

Article 22 (Transparency): Providing “meaningful information about the logic involved” for models with billions of parameters exhibiting emergent capabilities challenges current technical capabilities.¹⁵ The European Data Protection Board emphasizes that information must enable data subjects to understand and challenge decisions.¹⁶ Generic descriptions of “neural network architecture” fail this standard for high-risk contexts.

This requires organizations deploying open-weight models in high-risk contexts to somehow provide

12 Id. at p. 6: “Developers focused on technical issues like bias and safety, while researchers emphasized broader social impacts. Both groups paid little attention to fraud and manipulation, which are common harms arising from how people interact with AI.”

13 Hacker, P., Kilian, R., & Costas, J. (2025). “Simplifying” European AI Regulation: An Evidence-based White Paper. Bertelsmann Stiftung, p. 5 (quoting interview participant from SME legal tech sector).

14 Spanish Data Protection Agency (AEPD), “Adecuación al RGPD de tratamientos que incorporan Inteligencia Artificial” (2020), available at <https://www.aepd.es/>.

15 Regulation (EU) 2016/679 (General Data Protection Regulation), Article 22(3) (right to obtain human intervention, express one’s point of view and contest the decision requires “meaningful information about the logic involved”).

16 European Data Protection Board, Guidelines 8/2020 on the targeting of social media users (adopted 2 September 2020), para. 72-74.

meaningful explanations for systems they did not build and cannot fully audit, without the technical documentation, training data transparency, or provider cooperation that would make this feasible.

AI Act Compliance Challenge

The AI Act’s Article 25(2) requires original providers to “closely cooperate with new providers and shall make available the necessary information and provide the reasonably expected technical access and other assistance” when downstream modifications occur.¹⁷

However, general-purpose AI model providers are explicitly excluded from these collaboration duties.¹⁸ When an organization integrates an open-weight GPAI model into a high-risk application, it cannot compel cooperation from the original provider. This creates what regulatory experts call a “very significant regulatory gap”¹⁹, the Act encourages open-source adoption while providing no effective legal recourse for obtaining essential compliance information.

The SME Burden Multiplier

The compliance burden falls disproportionately on smaller organizations. One company managing over 4,000 AI applications described operator obligations as “an insane administrative effort.”²⁰ Interview evidence reveals that SMEs must hire both technical and regulatory personnel “from day one” for high-risk applications, representing “a challenge and a burden in terms of expense” that larger competitors more easily absorb.²¹

The resource asymmetry creates an unequal ecosystem, where large enterprises deploy dedicated legal teams and compliance infrastructure, while SMEs struggle to determine which obligations apply and how existing workflows relate to AI Act requirements.²²

17 AI Act, Article 25(2): “Providers of high-risk AI systems that continue to be providers of those systems after they have been placed on the market or put into service shall closely cooperate with new providers and shall make available the necessary information and provide the reasonably expected technical access and other assistance that are required for the fulfilment of the obligations set out in this Regulation...”

18 ocker & Holweg (2025), cited in Hacker et al. (2025), supra note 13, at p. 9: “However, GPAI model providers are excluded from Article 25(2), which creates a very significant regulatory gap in one of the most important industry use cases for AI adoption in the EU.”

19 Id.

20 ocker et al. (2025), supra note 13, at p. 6 (quoting interview participant from large industrial company: “The fact that we say I’m evaluating an AI system, but we see that in the future almost everything will be AI systems... This means for us that we also have corresponding obligations as operators there, and fulfilling these obligations – can you imagine what an insane administrative effort that is?”).

21 Id. at p. 6 (quoting Dr.-Ing. Julia Hoxha, CEO Zana Technologies GmbH: a start-up “has to hire both the CTO and the regulatory officer from day one, especially when you go into the space of being what they call a high-risk device... is a challenge and a burden in terms of expense”).

22 Id. at p. 4: “Companies report that determining which rules apply and how to demonstrate cross-regime compliance creates ‘an atmosphere of uncertainty and fear’, especially given that regulatory guidelines often arrive only shortly before legal provisions take effect, leaving minimal time for adaptation.”

“When an organization integrates an open-weight GPAI model into a high-risk application, it cannot compel cooperation from the original provider. This creates what regulatory experts call a “very significant regulatory gap”

This dynamic may inadvertently concentrate high-risk AI deployment among large providers, limiting the competitive and innovative benefits that open-weight models were intended to enable. As 96% of model card risk sections are exact duplicates²³ and regulatory guidance arrives shortly before legal provisions take effect, SMEs face an atmosphere of “uncertainty and fear.”²⁴

Conclusion

Organizations considering high-risk deployment of open-weight models should evaluate five factors:

1. **Genuine necessity:** Is this truly Annex III high-risk²⁵, or could lower-risk classification apply?
2. **Alternative architectures:** Proprietary API access maintains provider control and contractual compliance support. Structured access programs provide open research benefits with governance frameworks.
3. **Hybrid approaches:** Reserve open-weight models for lower-risk components; use fully controllable systems for high-risk decision points.
4. **Resource capacity:** Budget for substantial compliance infrastructure, documentation, audit capabilities, explainability tooling, monitoring systems, incident response, etc. Can the organization sustain this without provider support?
5. **Liability coverage:** Does professional liability insurance adequately cover algorithmic

23 Rao et al. (2025), supra note 8, at p. 3: “Among the 64,116 model cards [with risk sections], an overwhelming majority of risk sections (96%) were exact duplicates.”

24 Hacker et al. (2025), supra note 13, at p. 4 describing “atmosphere of fear and uncertainty” and “perceived time pressures”.

25 AI Act, Annex III (List of high-risk AI systems referred to in Article 6(2)).

decision-making risks, particularly given concentrated deployer liability?

For legal advisors, the threshold question remains, does the client have technical capacity and resources to meet full high-risk AI Act obligations without ongoing provider support? If not, open-weight deployment may present more risk than value.

While open-weight AI models offer genuine benefits, for high-risk applications under the AI Act, these benefits come with concentrated compliance obligations that many organizations, particularly SMEs, are unprepared to assume.

The regulatory framework’s risk-based approach is intentional, focusing scrutiny where fundamental rights are most at stake. But the structural information gap, Article 25(2) exclusion, and GDPR compliance challenges create a compliance trap, deployers assume full regulatory liability for systems they cannot fully understand, audit, or control.

Organizations must honestly assess whether they can navigate this terrain or whether proprietary alternatives with contractual compliance support better align with their capabilities and risk tolerance. The “free” in free and open-source AI may carry hidden regulatory costs that far exceed licensing fees.

About the Author:

Tanya Chib is the founder of Privacy Rules and specializes in AI governance, GDPR compliance, and EU AI Act implementation. With over a decade of legal experience, and expertise in healthcare AI, she advises clients ranging from Fortune 500 companies to medical device manufacturers on data protection, AI regulation, and commercial contracting.

