
Systematic Cross-Regulation Threat Topology for EU AI Governance¹

Wiktoria Leks
wiktoriazleks@gmail.com

Rían Czerwiński
i@rian.is

With
Apart Research

Abstract

A single frontier AI training run can simultaneously trigger obligations under the EU AI Act, GDPR, Copyright Directive, and NIS, yet no systematic framework maps these compounding regulatory threats across stakeholder types and jurisdictions. We present a systematic threat topology covering 19 EU-level regulations across five stakeholder categories (frontier model developers, deployment platforms, hardware providers, open-source developers, and research organizations), with geographic enforcement modifiers for all 27 Member States. Our methodology employs an activity-based stakeholder taxonomy, temporal activation mapping, and a KNOW/GUESS/UNKNOWN epistemic framework that quantifies regulatory uncertainty rather than obscuring it. Key findings include: (1) cross-regulation compounding creates multiplicative compliance surfaces where identical development activities trigger 3–5 regulatory regimes simultaneously; (2) enforcement concentration: five Data Protection Authorities (DPA) account for over 85% of €5.88B in cumulative GDPR fines, creates significant compliance cost differentials depending on establishment jurisdiction; (3) temporal cascading between February 2025 and August 2027 activates obligations under four major regulatory categories in overlapping waves, with August 2025 marking a critical inflection point for frontier AI providers. We release the full 41-page threat matrix as open infrastructure for practitioners navigating EU AI compliance during this implementation period.

¹ Research conducted at the [Technical AI Governance Challenge](#), 2026

1. Introduction

The EU’s regulatory approach to AI operates not through a single statute but through overlapping horizontal regulations, each with independent enforcement mechanisms, jurisdiction specific implementation, and distinct temporal activation schedules. A frontier lab training a foundation model on EU sourced data simultaneously faces obligations under the AI Act as a provider of general-purpose AI [1], GDPR as a data controller processing personal data [2], the Copyright Directive regarding text and data mining exceptions [3], and potentially NIS2 as a provider of essential services [4], each enforced by different authorities with different penalty structures. The AI Act alone imposes fines up to €35M or 7% of global turnover for prohibited practices; GDPR adds up to €20M or 4% of turnover; and these penalties compound independently across jurisdictions.

Current compliance guidance addresses these regulations in isolation. Law firm analyses cover individual statutes; DPA guidance is jurisdiction-specific; academic regulatory analysis operates at too high a level for practitioner use. This fragmentation creates two governance failures. First, practitioners cannot assess total compliance burden without synthesizing dozens of primary sources—a task requiring specialized legal expertise that disadvantages smaller organizations and safety-focused labs directing resources toward alignment rather than legal infrastructure. Second, cross-regulation interactions create compounding effects invisible to single-regulation analysis: a GDPR data subject access request may require disclosure of AI Act: mandated technical documentation, while Copyright Directive compliance may conflict with AI Act transparency obligations.

Our main contributions are:

- 1. A systematic threat taxonomy identifying cross-regulation compounding patterns across 19 EU regulations and 5 activity-defined stakeholder categories, revealing multiplicative compliance surfaces invisible to single regulation analysis.*
- 2. A geographic enforcement modifier framework mapping DPA enforcement patterns across all 27 Member States, quantifying order of magnitude compliance cost differentials arising from establishment jurisdiction.*
- 3. An epistemic classification methodology (KNOW/GUESS/UNKNOWN) that treats regulatory uncertainty as structured data, enabling practitioners to distinguish verified obligations from interpretive gaps requiring legal counsel.*

2. Related Work

EU AI regulatory analysis spans several domains. Policy research organizations have produced analyses of individual regulations, particularly the AI Act [5], providing valuable policy context but focusing on single statutes rather than cross-regulation interaction patterns. Major law firms have published AI Act implementation guides, but these are typically paywalled, single-regulation,

and oriented toward legal interpretation rather than systematic threat mapping. Model documentation frameworks: Model Cards [6], Datasheets for Datasets [7], and FactSheets [8] address transparency through structured disclosure but target individual documentation obligations rather than the full regulatory surface. Commercial compliance platforms offer regulatory management tooling but are proprietary, expensive, and typically scoped to data protection alone.

The gap our work addresses: no open-source, systematic, cross-stakeholder, geographically-aware threat topology exists for EU AI regulation. Existing resources are either single-regulation, single jurisdiction, paywalled, or too abstract for practitioner use. Our matrix provides the infrastructure layer, enabling systematic threat assessment across the complete regulatory surface.

3. Methods

Our methodology comprises five phases designed to maintain epistemic rigor, while producing practitioner usable output.

Source verification. *All regulatory text was sourced from EUR Lex official publications. Enforcement data was collected from DPA annual reports, EDPB registries, and the CMS GDPR Enforcement Tracker [9]. Each factual claim is tagged with epistemic status: KNOW (verified against primary legal source with article-level citation), GUESS (estimated from cited enforcement data with stated confidence basis), or UNKNOWN (critical implementation detail where authoritative guidance does not yet exist). This framework treats uncertainty as structured data, as a 2025 AI Act implementation detail marked UNKNOWN is more informative than its silent omission.*

Activity-based stakeholder taxonomy. *We classify stakeholders by regulatory-relevant activity rather than organizational type, yielding five categories: frontier model developers (training GPAI models above compute thresholds), deployment platforms (operating AI systems for end users), hardware/compute providers (manufacturing or supplying AI infrastructure), open-source developers (releasing model weights publicly), and research organizations. This captures regulatory exposure more accurately than firm level classification, since a single organization may occupy multiple categories simultaneously.*

Threat enumeration and cross-regulation analysis. *For each regulation–stakeholder intersection, we identify applicable articles, penalty structures, enforcement authority, and temporal activation date, producing approximately 95 primary threat cells ($19 \text{ regulations} \times 5 \text{ stakeholders}$). We then systematically map cases where a single development activity triggers obligations under multiple regulations simultaneously, identifying compounding patterns that amplify compliance burden beyond the sum of individual statutes.*

Geographic modifier framework. We profile all 27 Member State DPAs by enforcement history, staffing, budget, and notable decisions, classifying them into three enforcement tiers. These modifiers transform the regulation and stakeholder matrix from a binary compliance map into a jurisdiction sensitive threat surface, where identical activities carry substantially different risk profiles depending on establishment location.

4. Results

Systematic threat enumeration across 19 regulations and 5 stakeholder categories reveals several non-obvious structural patterns in the EU regulatory landscape.

Cross-regulation compounding. Table 1 summarizes threat severity across regulation stakeholder intersections. Frontier model developers and deployment platforms each face HIGH severity exposure under three or more independent regulatory regimes simultaneously. A single training run on EU sourced data triggers concurrent obligations under GDPR (lawful basis, Art. 6), the Copyright Directive (text and data mining, Art. 4), the AI Act (GPAI provider obligations, Art. 53), and potentially ePrivacy provisions: four independent compliance surfaces with four distinct enforcement authorities. Open source developers face a parallel compounding pattern, despite intuitions about reduced regulatory burden, AI Act GPAI provider obligations and Copyright Directive TDM requirements create HIGH-severity exposure comparable to proprietary frontier labs.

Table 1: Cross-Regulation Threat Severity Matrix. Levels reflect combined assessment of penalty magnitude, enforcement probability, and obligation complexity per regulation–stakeholder intersection.

Regulation	Frontier Labs	Platforms	Hardware	Open-Source	Research
GDPR	HIGH	HIGH	MEDIUM	LOW	MEDIUM
AI Act	HIGH	HIGH	MEDIUM	HIGH	LOW
Copyright Dir.	HIGH	MEDIUM	MINIMAL	HIGH	MEDIUM
DSA	LOW	HIGH	MINIMAL	LOW	MINIMAL
NIS2	MEDIUM	MEDIUM	HIGH	LOW	LOW
ePrivacy Dir.	MEDIUM	HIGH	LOW	LOW	LOW
DMA	LOW	HIGH	MINIMAL	LOW	MINIMAL
DORA	LOW	MEDIUM	MEDIUM	MINIMAL	MINIMAL

Enforcement concentration and jurisdictional arbitrage. GDPR enforcement data reveals extreme concentration: cumulative fines since 2018 total €5.88B across 2,245 recorded penalties [9, 10]. Ireland’s DPC alone accounts for €3.5B (60%), driven by its one-stop-shop status as lead supervisory authority for technology companies headquartered in Dublin. The top five DPAs by fine value (Ireland, Luxembourg, France, Italy, the Netherlands) account for over 85% of total penalty volume. This creates measurable jurisdictional arbitrage: identical GDPR regulated activities carry fundamentally different enforcement risk depending on the controller’s EU establishment,

with the Ireland–Poland differential representing roughly an order-of-magnitude difference in expected penalty exposure.

Temporal activation cascading. The EU AI regulatory landscape activates in overlapping waves (Table 2). GDPR and DSA are already fully enforced, while AI Act obligations phase in across four deadlines between February 2025 and August 2027. The August 2025 inflection point is critical: GPAI provider obligations activate, requiring frontier labs to publish model documentation, implement copyright compliance procedures, and conduct systemic risk assessments, all while existing GDPR and Copyright Directive obligations remain independently enforced.

Table 2: Temporal Activation Cascade. Phased regulatory activation creates compounding compliance deadlines, with August 2025 marking a critical inflection for frontier AI providers.

Date	Regulation	Key Provisions	Primary Stakeholders
Active (2018–)	GDPR	Full enforcement; €5.88B cumulative fines	All
Active (2022–)	DSA / DMA	Platform transparency, gatekeeper obligations	Platforms
Oct 2024	NIS2	Transposition deadline; cybersecurity obligations	Hardware, Platforms
Feb 2025	AI Act Ch. I–II	Prohibited practices ban; AI literacy	All
Aug 2025	AI Act Ch. III, V	GPAI transparency; systemic risk assessment	Frontier Labs, Open-Source
Aug 2026	AI Act Ch. III §2	High-risk system conformity assessment	Platforms, Hardware
Aug 2027	AI Act (full)	Complete enforcement; remaining provisions	All

Epistemic gaps as structured finding. Application of the KNOW/GUESS/UNKNOWN framework reveals that a substantial proportion of practitioner critical implementation details remain formally unresolved. Key UNKNOWN areas include: the EU AI Office’s interpretation of GPAI “systemic risk” thresholds, the scope of open source exemptions under AI Act Art. 2(12), interaction between GDPR legitimate interest and AI Act transparency requirements, and Member State transposition choices for NIS2 penalty regimes. These gaps are not omissions they are findings about the current state of regulatory certainty.

5. Discussion and Limitations

These structural patterns carry implications beyond individual compliance decisions.

Cross-regulation compounding disproportionately burdens smaller organizations, since compliance costs scale sub-linearly with organizational size, large incumbents with dedicated legal teams face proportionally lower per-model costs than startups directing resources toward alignment.

Jurisdictional enforcement arbitrage creates incentives to establish in low enforcement Member States, a race to bottom dynamic undermining the EU’s harmonization objectives. The 2025–2027 window represents a period of acute regulatory risk where single regulation compliance strategies will prove insufficient.

5.1 Limitations

Several constraints bound these findings. GUESS estimates lack project-specific validation and should be treated as baselines requiring organizational calibration. UNKNOWN gaps reflect genuine epistemic boundaries, particularly around AI Act implementation where EU AI Office guidance remains forthcoming. The matrix is a temporal snapshot as of February 2026.

Geographic analysis draws primarily on GDPR enforcement data, as AI Act enforcement has not yet generated comparable patterns. Scope is EU only; US, UK, and Chinese regulatory frameworks create additional surfaces not addressed here.

5.2 Future Work

Natural extensions include dynamic updating infrastructure as regulations amend and enforcement decisions accumulate; jurisdiction expansion to US, UK, and Chinese frameworks; automated EUR Lex and DPA registry monitoring; and case law integration as AI Act enforcement generates judicial interpretation.

6. Conclusion

We present a systematic threat topology mapping 19 EU regulations across 5 activity-defined stakeholder categories with geographic enforcement modifiers for all 27 Member States. The EU AI regulatory landscape operates as an interacting system: cross-regulation compounding creates multiplicative compliance surfaces, enforcement concentration enables jurisdictional arbitrage, and temporal cascading produces compliance cliff effects during the 2025–2027 window. The KNOW/GUESS/UNKNOWN framework quantifies not only what is regulated but what remains uncertain and critical for practitioners under genuine ambiguity. The accompanying 41-page matrix provides open infrastructure extensible to additional jurisdictions as enforcement patterns emerge.

Code and Data

Primary artifact: 41-page [EU AI Regulatory Threat Matrix \(PDF\)](#), containing complete threat enumeration across all regulation stakeholder intersections with article level citations, DPA enforcement profiles for all 27 Member States, and temporal activation mapping (<https://github.com/rianczerwinski/EU-AI-regulatory-threat-matrix/blob/main/EU-AI-regulatory-threat-matrix.tex>)

References

1. *European Parliament and Council of the EU. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). Official Journal of the European Union, 2024.*
2. *European Parliament and Council of the EU. Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (GDPR). Official Journal of the European Union, 2016.*
3. *European Parliament and Council of the EU. Directive (EU) 2019/790 on copyright and related rights in the Digital Single Market. Official Journal of the European Union, 2019.*
4. *European Parliament and Council of the EU. Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2). Official Journal of the European Union, 2022.*
5. *Renda, A. et al. “Study to support an impact assessment of regulatory requirements for AI in Europe.” European Commission, 2021.*
6. *Mitchell, M. et al. “Model Cards for Model Reporting.” Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*), 2019.*
7. *Geburu, T. et al. “Datasheets for Datasets.” Communications of the ACM 64(12), 2021.*
8. *Arnold, M. et al. “FactSheets: Increasing Trust in AI Services through Supplier’s Declarations of Conformity.” IBM Journal of Research and Development 63(4/5), 2019.*
9. *CMS Law. GDPR Enforcement Tracker Report 2024/2025. CMS, 2025.*
<https://www.enforcementtracker.com/>
10. *DLA Piper. GDPR Fines and Data Breach Survey: January 2025. DLA Piper, 2025.*
11. *European Parliament and Council of the EU. Regulation (EU) 2022/2065 on a Single Market for Digital Services (DSA). Official Journal of the European Union, 2022.*

LLM Usage Statement

Claude (Anthropic) was used extensively for regulation text synthesis, cross-regulation interaction analysis, threat classification logic, and document drafting. All legal citations were verified against EUR Lex primary sources. All DPA enforcement data was verified against published registries and the CMS GDPR Enforcement Tracker [9]. GUESS estimates are sourced from cited enforcement data and industry studies. UNKNOWN gaps are explicitly marked where authoritative guidance does not yet exist. The KNOW/GUESS/UNKNOWN framework itself functions as the verification methodology: LLM outputs are systematically categorized by epistemic status, ensuring that unverified claims are surfaced rather than silently incorporated.