POST /api/v1/compliance/init

Privacy & AI Law — S06E06: Personal Data in AI Products

Training, Prompts, Automated Decisions & Vendor DPAs

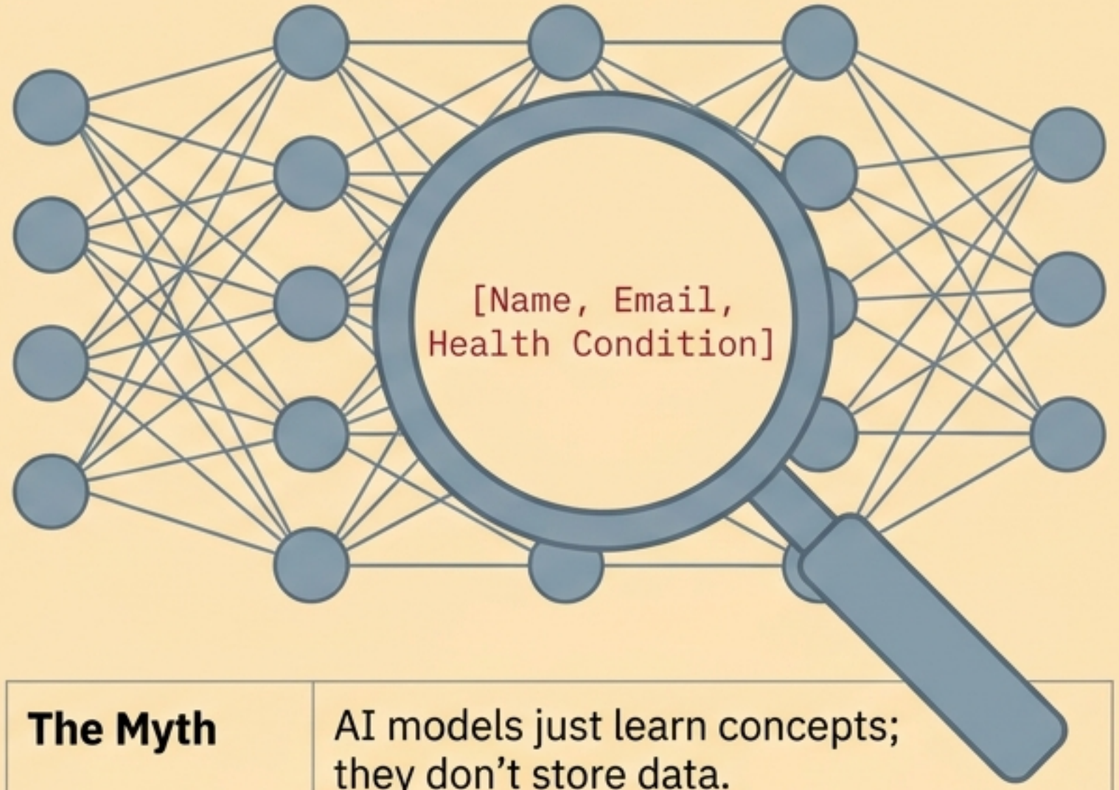
🛠️ For Founders, Product Managers, Developers, and Marketers

⚠️ General information, not legal advice — current as of November 2026.

Gemini Notebook

The Legal Hierarchy

S06E05 Recap: The EU AI Act imposes a tiered risk system—but if an AI touches personal data, the GDPR applies first.

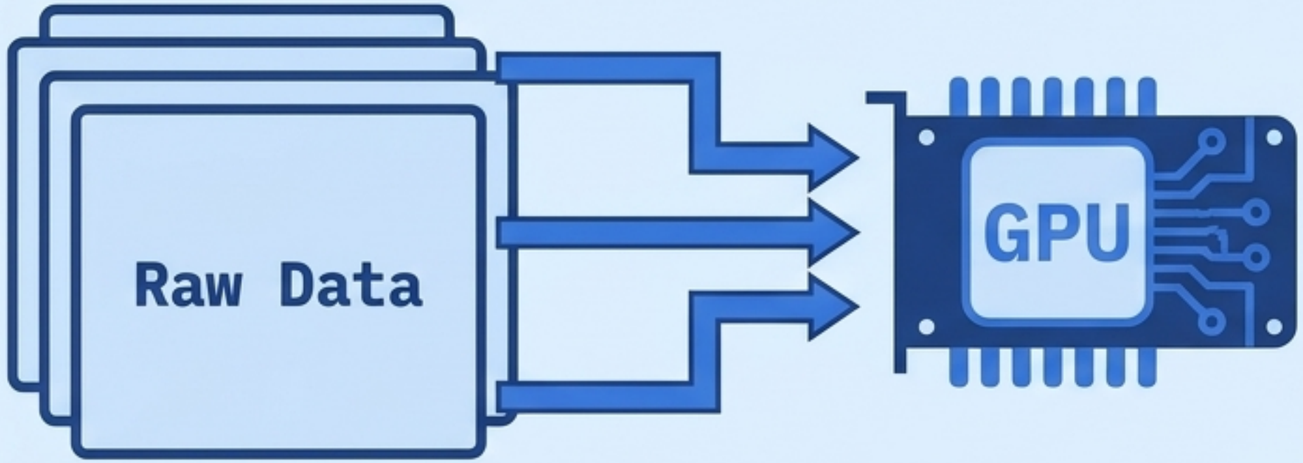


The Myth	AI models just learn concepts; they don’t store data.
The Reality	EDPB Opinion 28/2024 & NIST AI 600-1 confirm “**Data Memorization.**” LLMs can leak, generate, or correctly infer PII from their training data.

The Rule: If a model’s weights can extract or reconstruct an individual’s personal data, the model itself is processing personal data.

Gemini Notebook

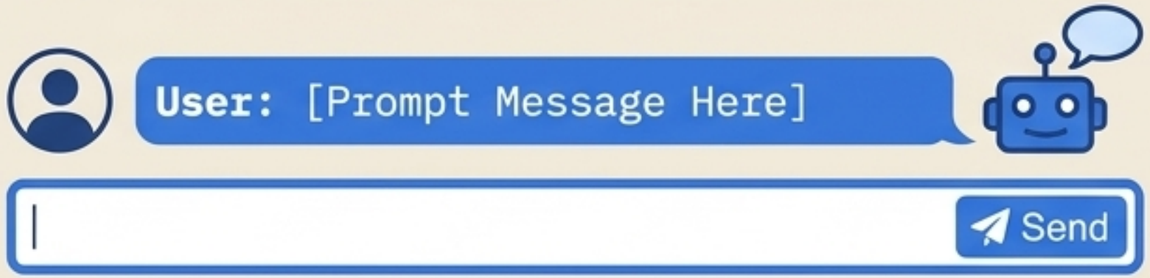
Phase 1: Development (Training & Fine-Tuning)



UK ICO & CNIL

Requires its own distinct Lawful Basis (e.g., Legitimate Interest).

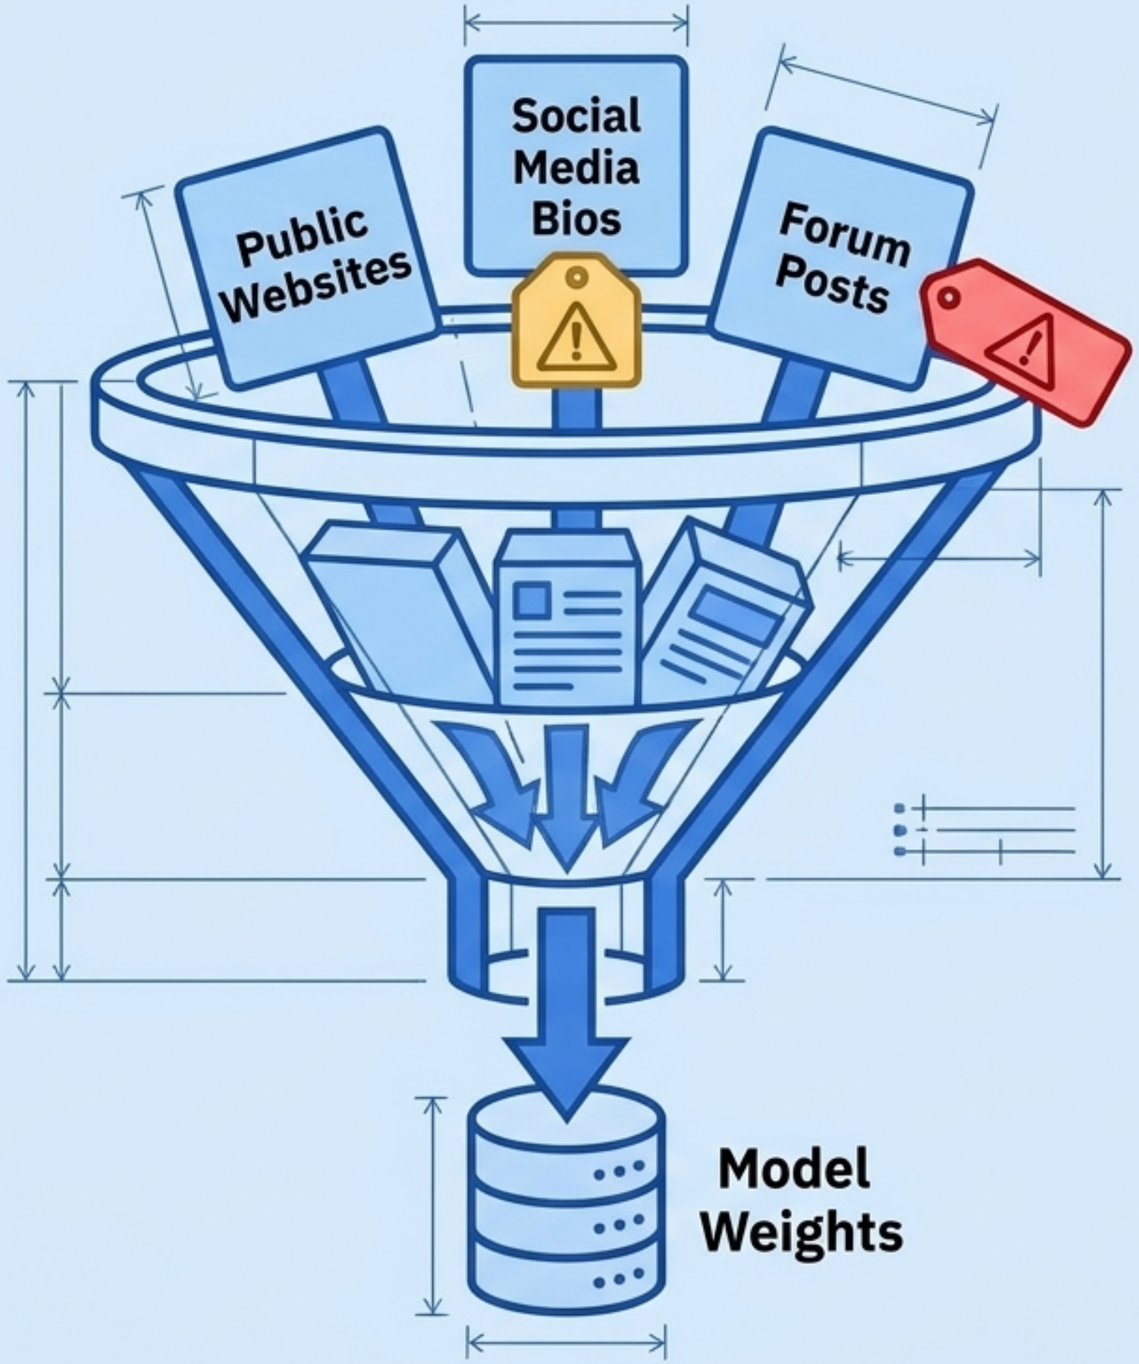
Phase 2: Deployment (Inference & Prompting)



Requires an entirely separate Lawful Basis (e.g., Contractual Necessity or Consent).

Key Takeaway: You cannot rely on a user's deployment consent to justify the upstream training of the model.

Gemini Notebook



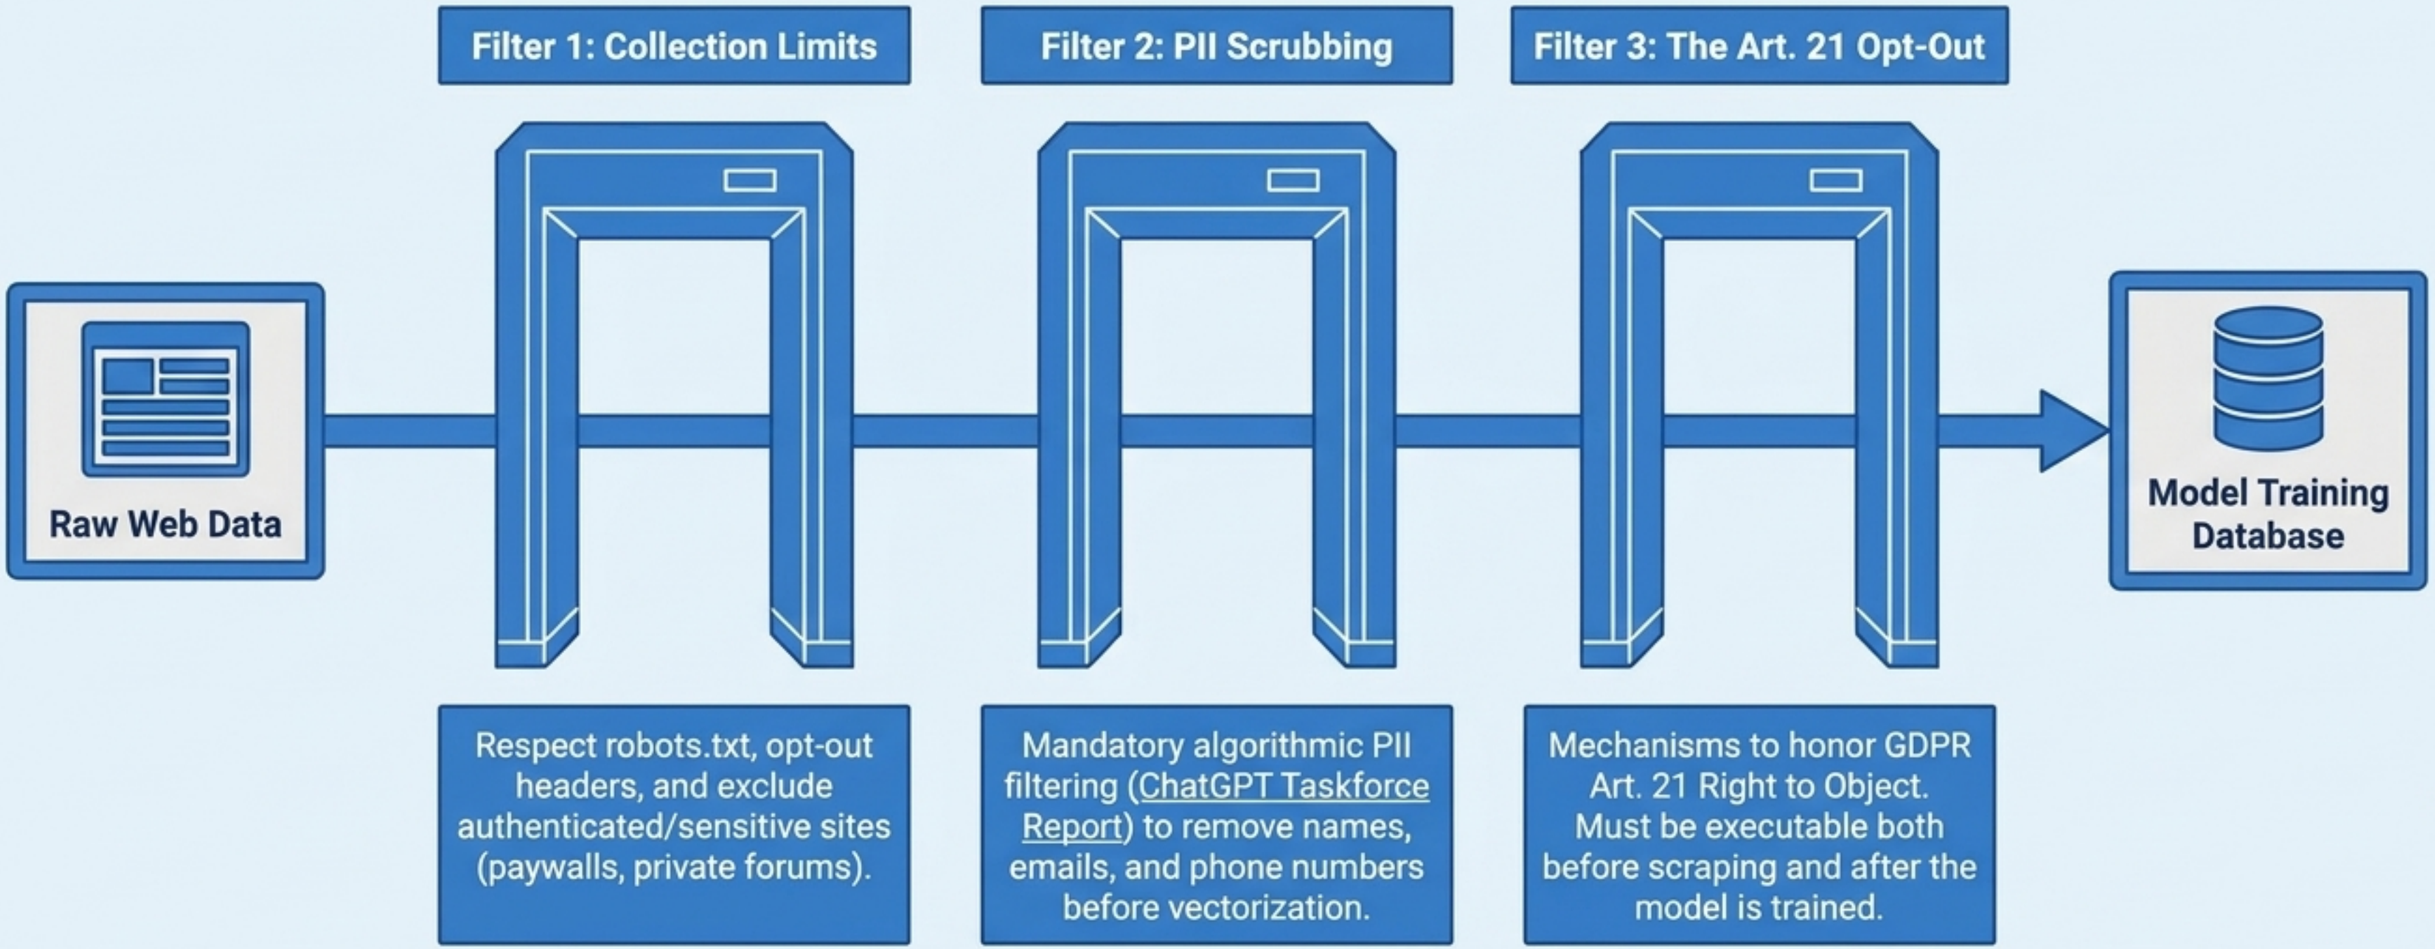
The Web-Scraping Minefield

GDPR Art. 6(1)(f) - Legitimate Interest

- EDPB 1/2024 & CNIL GenAI Sheets confirm that scraping public data for AI training can rely on Legitimate Interest.
- However, it is highly conditional and subject to strict balancing tests.
- Core Principle: ‘Publicly accessible’ does NOT mean ‘exempt from GDPR.’

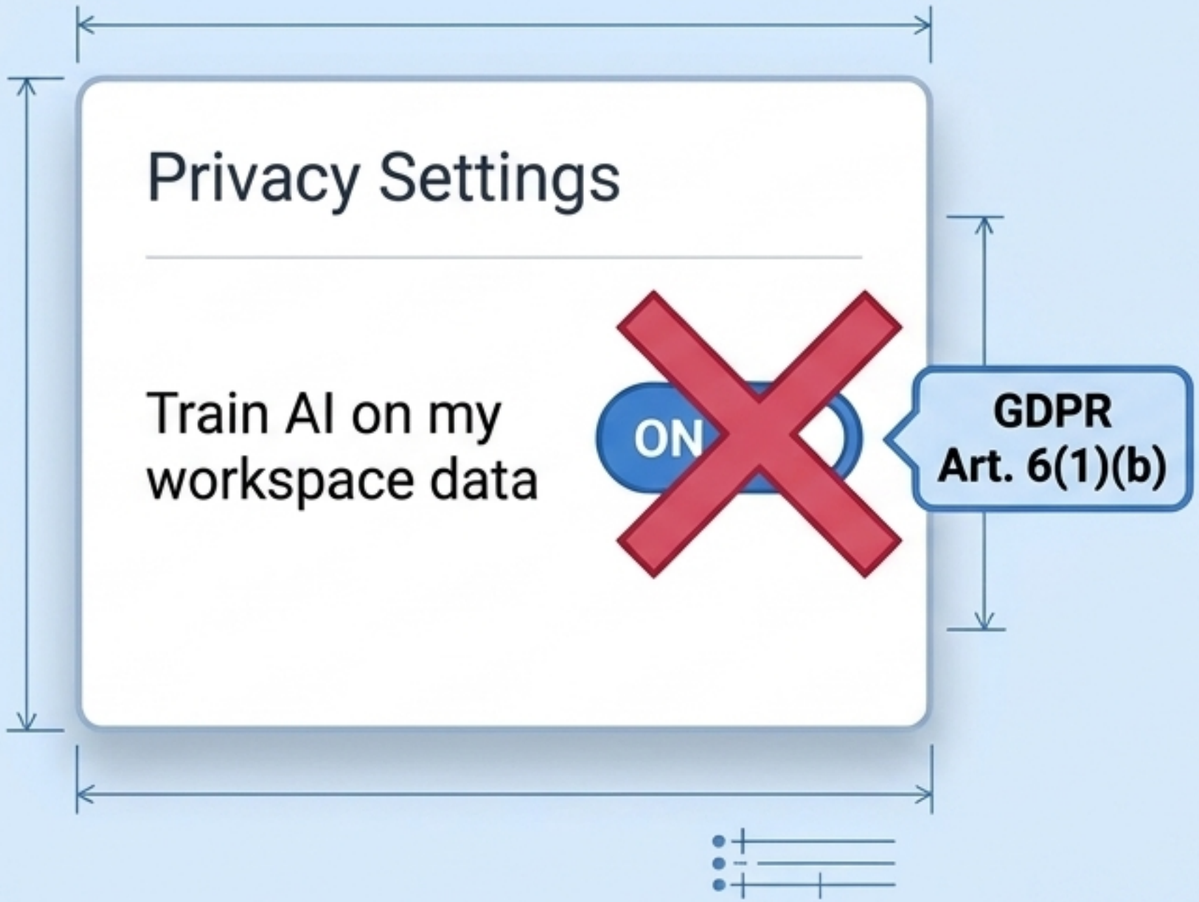
Gemini Notebook

Art. 6(1)(f) Scraping Safeguards & Opt-Outs



Gemini Notebook

The ‘Default-On’ Trap for Customer Data



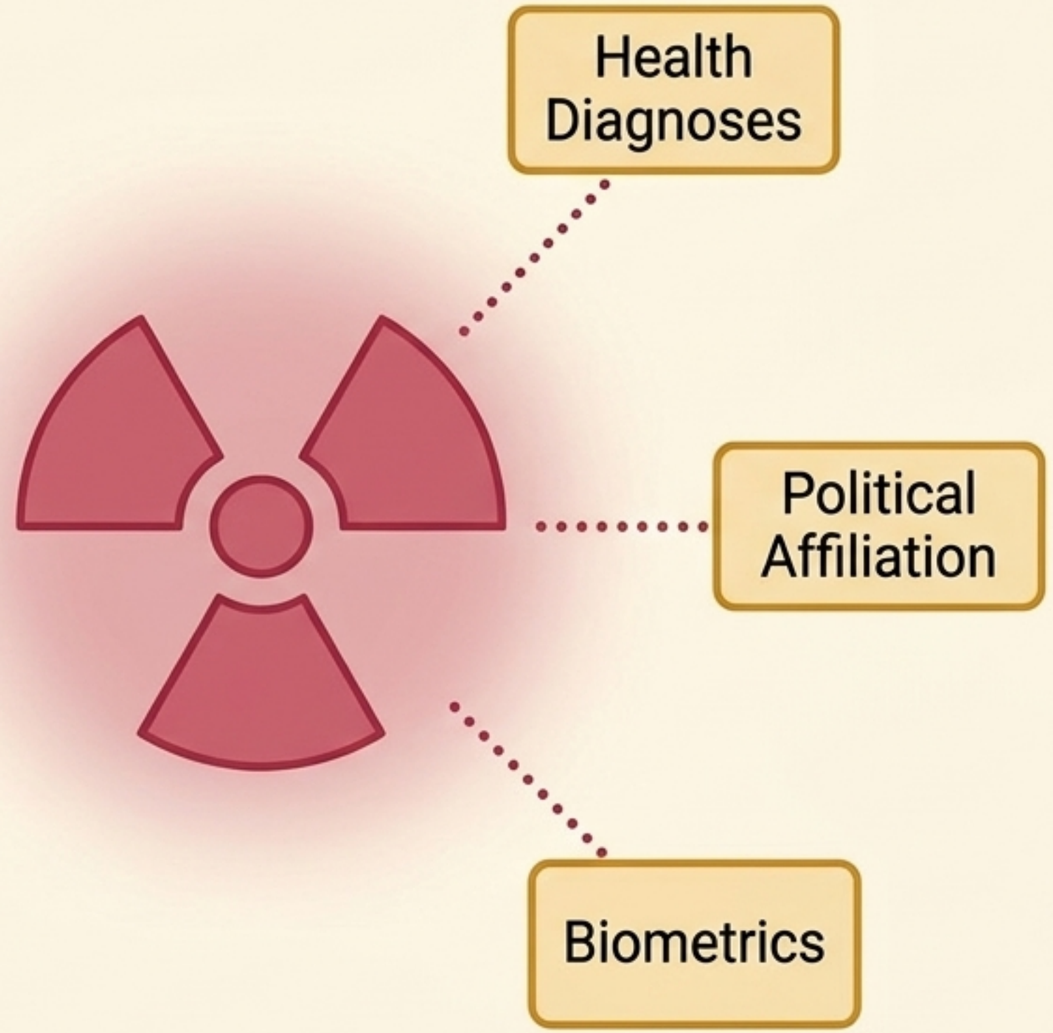
The Trap: You cannot hide customer-data training in your Terms of Service under Art. 6(1)(b) Contractual Necessity.

The Precedent: EDPB 05/2020 dictates that “improving our services via AI” is not strictly necessary to provide the core SaaS contract.

Regulatory Reality: The FTC is aggressively pursuing retroactive-ToS enforcement for bait-and-switch data training. Training requires explicit, opt-in Consent or a flawlessly executed Legitimate Interest assessment.

Gemini Notebook

Accidental Art. 9 Special-Category Ingestion



The Art. 9 Trap: Accidental ingestion of GDPR Art. 9 Special Category Data is strictly prohibited in standard trainig without explicit consent or substantial public interest.

The AI Act Exception (Art. 10(5))
The EU AI Act introduces a strict, narrow exception. Providers may exceptionally process Art. 9 data only if strictly necessary for “bias detection and correction.” Requires state-of-the-art security, strict pseudonymization, and immediate deletion once the bias is corrected.

Gemini Notebook

AI Architecture Privacy Comparison Table

	1. Base-Model Fine-Tuning (Updating Weights)	2. Access-Controlled RAG (Vector DB + API)	3. Zero-Retention LLM API (Pass-through)
Lawful Basis Friction	**High** (Requires re-training justification)	Medium (DB storage)	**Low** (Transient)
Cross-Tenant Leakage Risk	**High** (Weights can memorize)	**Low** (Hard logic access controls)	**None**
Hallucination Controls (Art. 5(1)(d) Accuracy)	Probabilistic	Deterministic grounding	Probabilistic

(Matrix dimensions grounded in CNIL GenAI Q&A and NIST AI 600-1)

Gemini Notebook

Fulfilling Data Subject Rights across Architectures (Art. 15, 16, 17, 21)

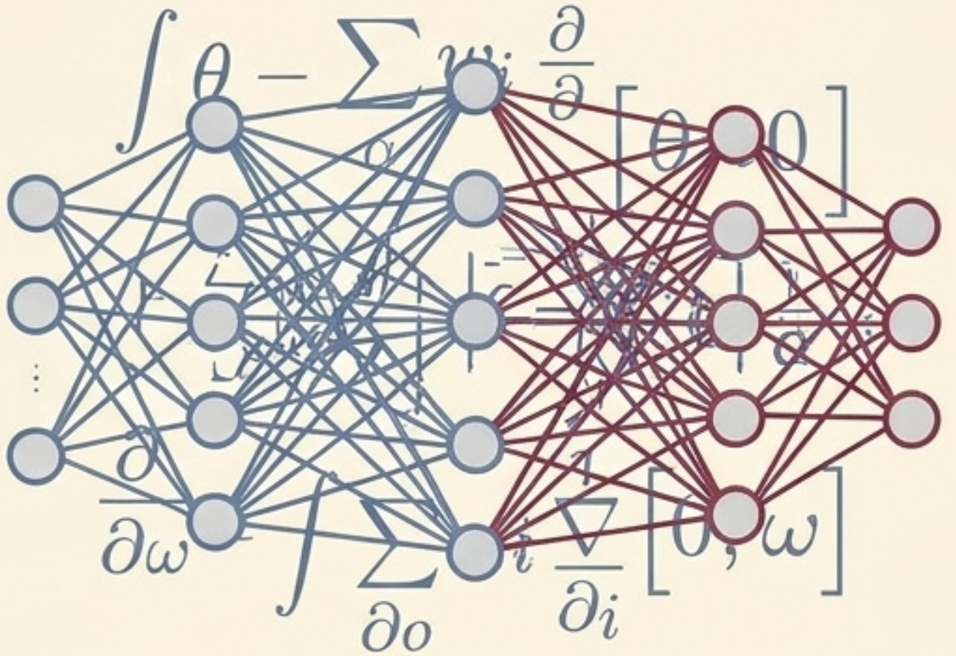
RAG Architecture

```
DELETE FROM vector_db
WHERE user_id = '123'
```



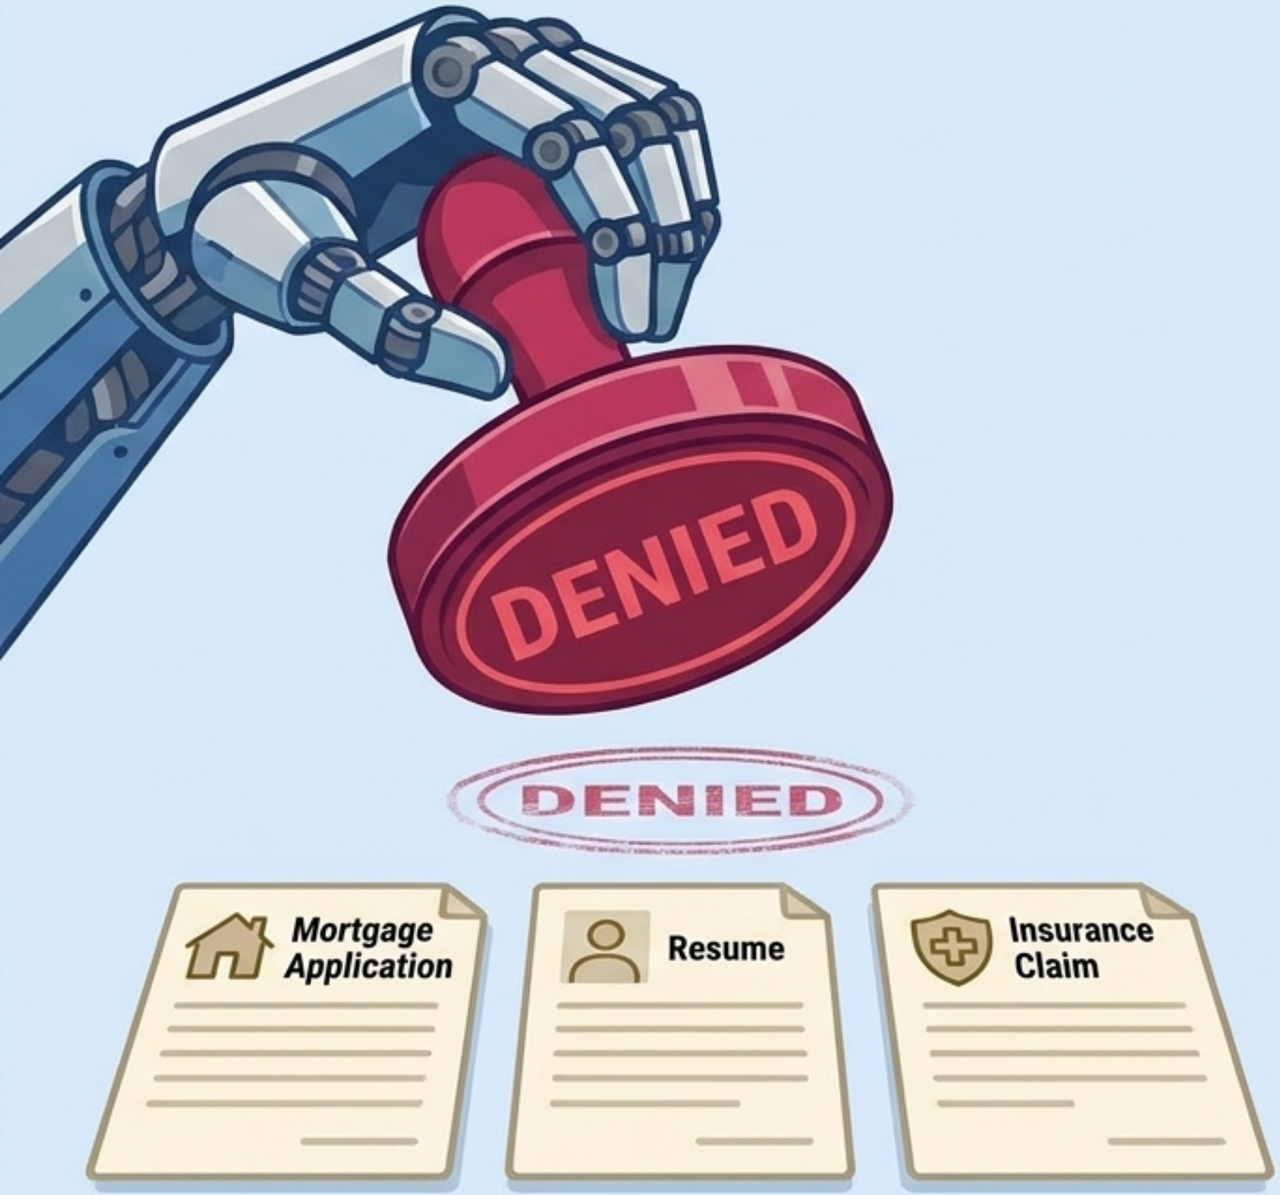
Trivial. If a user requests Art. 17 Erasure, you delete their document from the Vector DB. The AI instantly stops retrieving it.

Fine-Tuning Architecture



Near-Impossible. How do you unlearn a specific person's data from a billion-parameter model's weights without retraining the entire foundation model from scratch?

Gemini Notebook



Art. 22 & Solely Automated Decisions

GDPR Art. 22

EU AI Act Art. 26

The Rule:
Individuals have the right not to be subject to a decision based solely on automated processing which produces legal or similarly significant effects.

The Scope:
Regulators heavily scrutinize credit scoring, algorithmic hiring, insurance underwriting, and automated account bans.

Gemini Notebook

The Landmark CJEU SCHUFA Ruling (Case C-634/21)



The Precedent: If an AI calculates a credit score, and a human determinatively "rubber-stamps" that score to make a final decision, the creation of the score itself constitutes Automated Decision-Making under Art. 22.

The Lesson: Performative human intervention does not save you. WP251 rev.01 dictates the human must have meaningful authority and alternative data to override the AI.

Gemini Notebook

Exceptions, Safeguards & DPIAs

GDPR Art. 22(2) GDPR Art. 35

Exception Check
Legal automated decisions must be based on Explicit Consent, Contractual Necessity, or Union/Member State Law.

Mandatory Safeguards
Deployers must provide meaningful information about the logic involved, the right to obtain human intervention, and the right to contest the decision.

Mandatory DPIA
An Art. 35 Data Protection Impact Assessment is non-negotiable before deployment.

(Note: The California CPPA heavily mirrors these ADMT transparency rights).

Gemini Notebook

Point-of-Prompt vs. Scraped-Data Notices

Direct Collection (Art. 13)



Point-of-Prompt Notices:

Point-of-Prompt Notices: When collecting data directly users via prompts, warn them at the point of input not to share unnecessary third-party or Art. 9 PII.

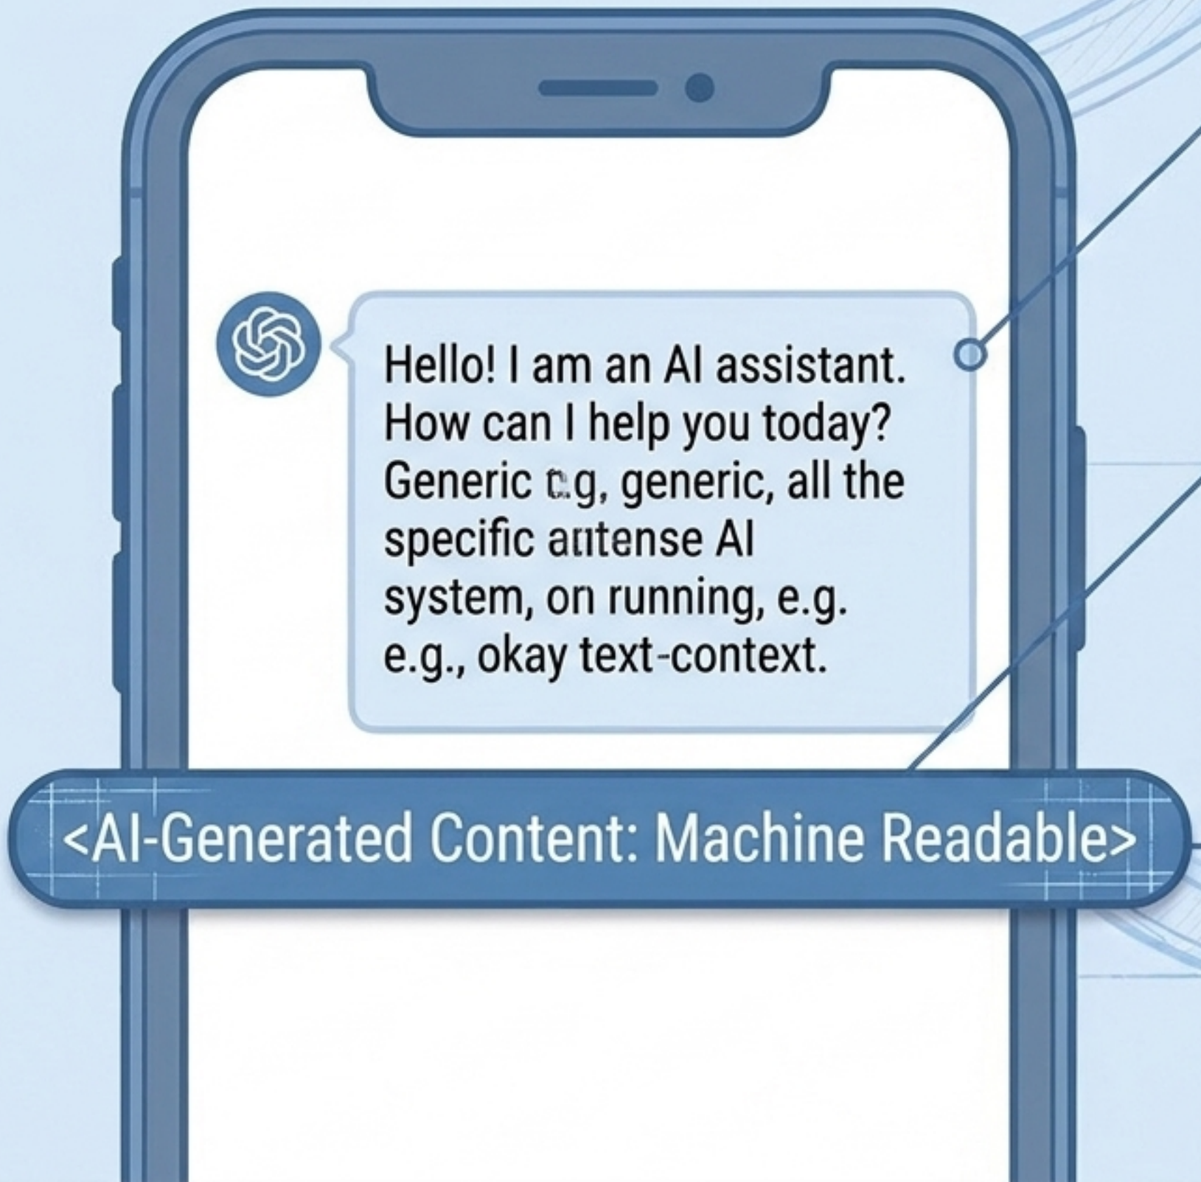
Indirect Collection (Art. 14)



Scraped-Data Notices: When training on scraped data, contacting every individual is impossible. You must rely on the Art. 14(5)(b) Disproportionate Effort exemption, requiring robust public privacy notices detailing exactly what data categories are scraped.

Gemini Notebook

EU AI Act Art. 50 Chatbot & Deepfake Disclosures



Chatbots (Art. 50(1)): Users must be explicitly informed they are interacting with an AI system.

Deepfakes (Art. 50(2)): Synthetic audio, video, and text must be marked in a machine-readable format and detectable as artificially generated.

Hallucination Warnings (GDPR Art. 5(1)(d)): Deployers must warn users that outputs may confabulate facts about real people.

Gemini Notebook

Mandatory Art. 28 Processor DPA Clauses



The Non-Negotiable Clause

Your LLM API vendor (e.g., OpenAI, Anthropic) must act legally as a Processor, executing your documented instructions (Art. 28(3)(a)).

The DPA must explicitly block the vendor from using your API inputs and outputs for their own foundation model training.

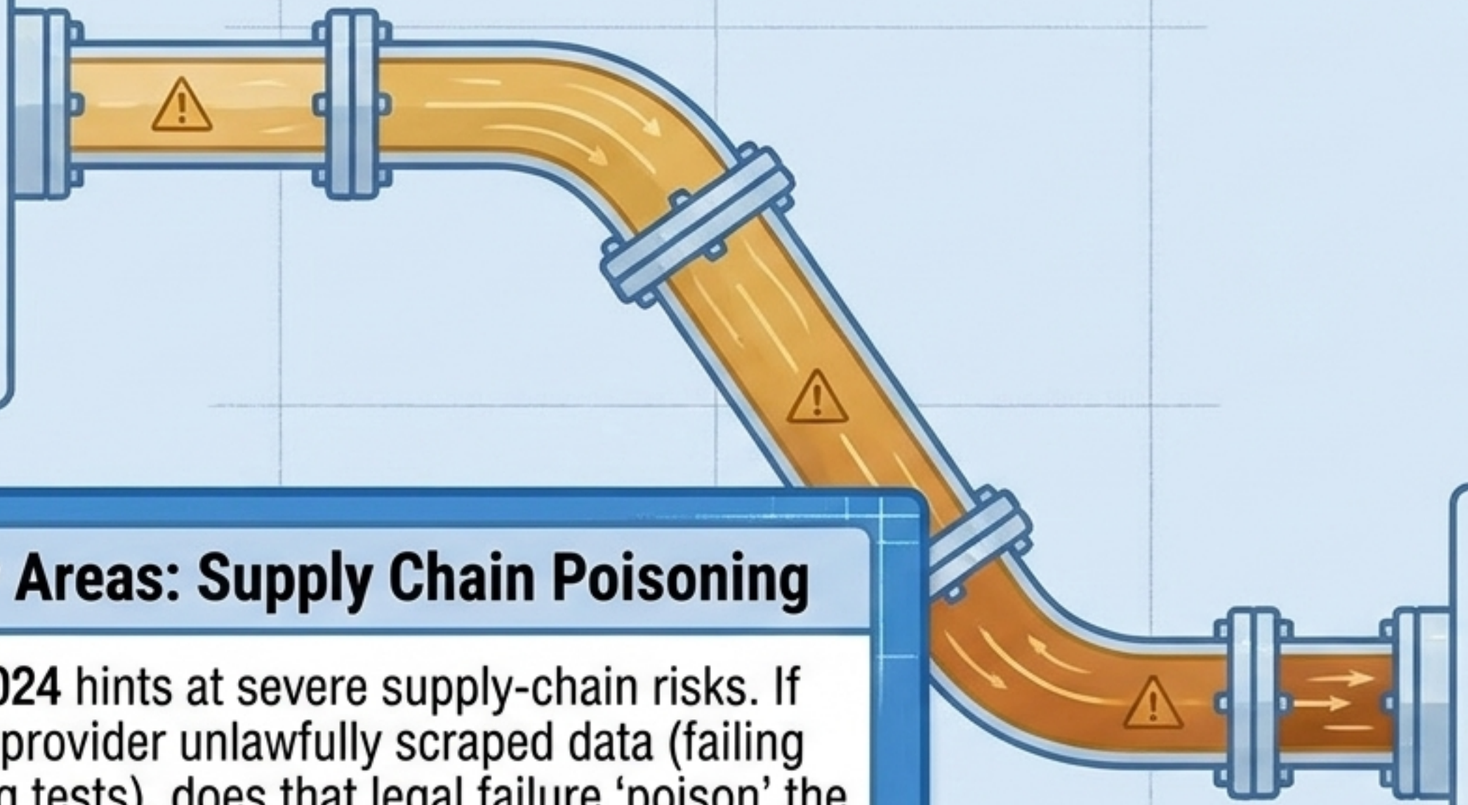
Sub-processors: Ensure flow-down clauses apply to any sub-vendors the API uses (Decision 2021/915 & EDPB Opinion 22/2024).

Gemini Notebook

The Downstream Deployer Dilemma



Foundation Model Provider





Your Product (Deployer)

Contested Grey Areas: Supply Chain Poisoning

EDPB **Opinion 28/2024** hints at severe supply-chain risks. If a foundation model provider unlawfully scraped data (failing Art. 6(1)(f) balancing tests), does that legal failure 'poison' the weights? Deployers face active regulatory scrutiny over whether using a legally compromised model constitutes an ongoing processing violation in their own downstream product.

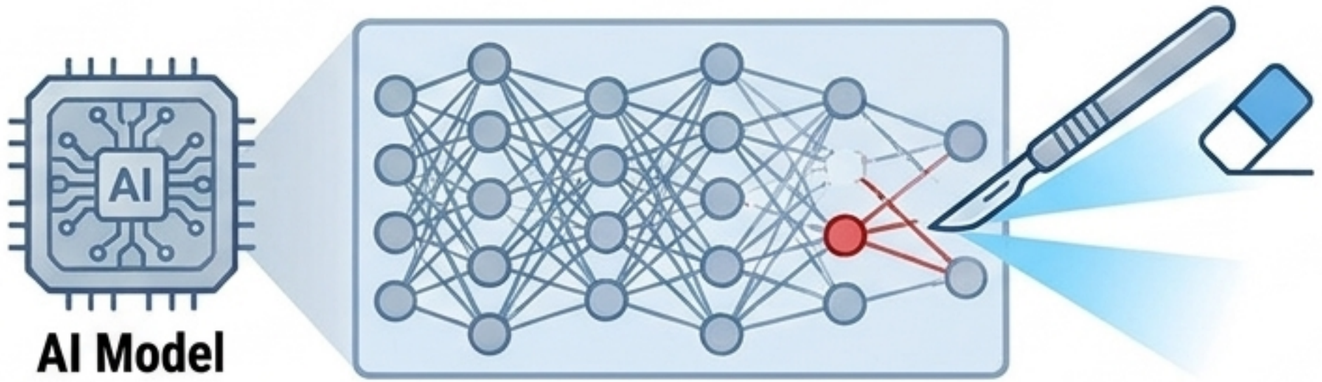
Gemini Notebook

Machine Unlearning vs. Output Filters (Art. 17 Erasure)

Method 1: Output Filter (Industry Standard)



Method 2: Machine Unlearning (Active Research)



The Regulatory Debate:

Does an output filter that merely prevents the model from saying a deleted user's name satisfy the Right to Be Forgotten?

Or must the deployer **mathematically extract** the data from the weights?

(EDPB ChatGPT Taskforce vs. ICO GenAI response).

Gemini Notebook

Probabilistic Hallucinations of Real People

John Doe filed for ~~bankruptcy~~ in ~~2018~~ and served three years for ~~fraud~~.

GDPR Art. 5(1)(d)

The Regulatory Stance:

- **The Accuracy Principle**: GDPR requires personal data to be “accurate and, where necessary, kept up to date.”
- **The Technical Reality**: LLMs do not store databases; they predict tokens. NIST AI 600-1 defines this as “confabulation.”
- **The Regulatory Stance:** When an LLM hallucinates false facts (criminal records, financial ruin) about a real person, regulators view this as a **strict accuracy violation**, requiring immediate correction mechanisms.

Gemini Notebook

Five Engineering Defaults for Shipping Safely



1 Zero-Retention APIs

Default to enterprise API tiers with explicit “Zero Data Retention” DPAs.



2 RAG over Fine-Tuning

Use Vector DBs to ground data. Avoid baking user data into model weights.



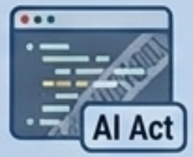
3 Opt-In Toggles

Never default-on user data training. Require explicit opt-in.



4 Point-of-Prompt Warnings

Build inline UI reminders instructing users not to paste PII.

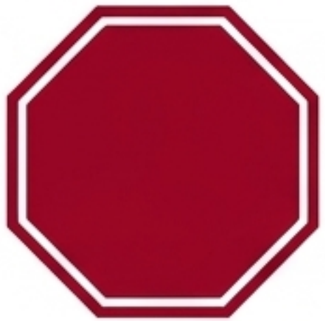


5 AI Watermarking

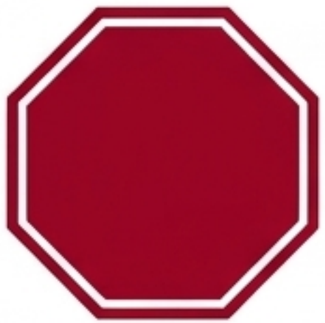
Comply with AI Act Art. 50 by marking all synthetic generations.

Gemini Notebook

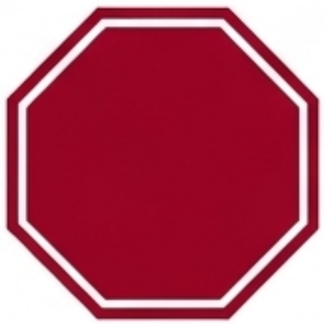
Counsel Red Lines (Stop & Escalate)



Training on Customer/Scraped PII: Attempting to fine-tune or train models on personal data without explicit consent or a documented Legitimate Interest Assessment.



Art. 22 ADMT Scoring: Building systems that automatically approve or deny housing, employment, credit, or vital services without meaningful human override.



Art. 9 Prompts: Actively prompting for, ingesting, or inferring health data, biometrics, or political affiliations without explicit user consent.

Gemini Notebook

POST /api/v1/compliance/complete

Conclusion

Privacy by Design must be embedded at the architectural level—from the choice of model to the configuration of the DPA.

Next Up

→ S06E07: A Practical Checklist (The Season Finale)

⚠ General information, not legal advice — current as of November 2026.

Gemini Notebook