The EU AI Act in Plain Language

FREE EDUCATIONAL BRIEF · LEARN AI. UNCOVER THE DATA. PROTECT THE TRUTH.

The EU AI Act is the first comprehensive, binding legal framework for artificial intelligence — and because global companies rarely build separate products per region, its requirements travel far beyond Europe. The core design is simple to state: regulate by risk, not by technology. The riskier the use, the heavier the obligations.

THE RISK PYRAMID

» **Unacceptable risk — banned**

Social scoring by governments, manipulative techniques causing harm, untargeted scraping of faces for recognition databases, emotion inference in workplaces and schools (with narrow exceptions).

» **High risk — strictly managed**

AI in hiring, credit, education, critical infrastructure, medical devices, law enforcement: requires risk management, quality data, documentation, human oversight, and registration.

» **Limited risk — transparency**

Chatbots must identify as machines; synthetic media and deepfakes must be labeled as such.

» **Minimal risk — untouched**

Spam filters, game AI, recommendation helpers: no new obligations.

GENERAL-PURPOSE MODELS GET THEIR OWN CHAPTER

Providers of large general-purpose models owe technical documentation, a summary of training content, and copyright-policy compliance; models above a compute threshold deemed “systemic risk” face added testing, incident reporting, and cybersecurity duties. Penalties scale to global revenue — designed to matter even to the largest labs.

WHAT IT MEANS IN PRACTICE

If you deploy AI toward EU users, classify your use case first; most business automation lands in minimal-or-limited risk, where the main duty is honesty — disclose bots, label synthetic media. The Act's transparency defaults align exactly with provenance practice: label what is generated, document what your systems do, keep humans over consequential decisions. Teams doing that today are largely compliant by design.

INVESTIGATOR'S TAKEAWAYS

- Regulation follows the use case's risk, not the model's brand.
- A short list of practices is banned outright.
- High-risk uses need documentation, oversight, and registration.
- Bots must disclose; deepfakes must be labeled.
- Provenance-first workflows are compliance-first workflows.