AI LITERACY SERIES

The EU AI Act: A Plain-English Guide

Everything you need to know before reading the detail

Ryan Bishop | March 2026 | ryanbishop.co.uk

1. What is the EU AI Act?

The EU AI Act (Regulation EU 2024/1689) is the world's first comprehensive legal framework for artificial intelligence. Published in the Official Journal of the European Union on 12 July 2024 and entering force on 1 August 2024, it applies across all 27 EU member states without needing separate national legislation.

Its stated purpose is straightforward: ensure AI used in the EU is safe, transparent, and respects fundamental rights — while keeping the door open for innovation. In practice, it achieves this by classifying AI systems into risk tiers and attaching different obligations to each.

One important distinction: the Act does not regulate AI as a technology. It regulates **AI systems** — specific applications deployed for specific purposes. The same underlying model can be minimal risk in one deployment and high risk in another, depending on what it is used to do.

2. Who does it apply to?

The Act has broad geographic reach. It applies to anyone who:

- Places an AI system on the EU market — regardless of where they are based
- Puts an AI system into service within the EU
- Uses an AI system where its outputs affect people in the EU

This means a company headquartered in the US, UK, or anywhere else is still subject to the Act if its AI systems operate in the EU.

The provider / deployer distinction

The Act assigns different obligations depending on your role:

Role	Definition	Obligations
Provider	Develops and places an AI system on the market or puts it into service — under their own name. This includes building on top of an existing model.	Heaviest. Technical documentation, conformity assessment, registration, CE marking (high-risk systems).
Deployer	Uses an AI system developed by another party for a specific professional purpose.	Lighter, but real: human oversight, transparency to those affected, logging, incident reporting for high-risk systems.

Key point: you can be both a provider and a deployer simultaneously — for different AI systems, or even within the same engagement.

3. The Four Risk Tiers

The Act classifies every AI system into one of four tiers. The tier determines your obligations.

Risk Tier	What it covers	Who it affects	Obligations
Unacceptable / Prohibited	Social scoring, subliminal manipulation, most real-time biometric surveillance	Providers & deployers	Complete ban — cannot be placed on market
High Risk	AI in medical devices, employment decisions, education access, law enforcement, critical infrastructure, biometric ID	Providers primarily; deployers have oversight duties	Full Chapter III: technical docs, conformity assessment, registration, human oversight, logging
Transparency	Chatbots, deepfake/synthetic content tools, emotion recognition	Providers & deployers	Must declare AI interaction; watermark synthetic content
Minimal Risk	Everything else — spam filters, AI in video games, recommendation engines, productivity tools	Providers & deployers	No mandatory obligations — voluntary codes of conduct encouraged

How classification works: there are two routes to high-risk status. **Route 1 (Annex I):** the AI is a safety component of a product already subject to EU harmonisation legislation (medical devices, machinery, aviation, etc.). **Route 2 (Annex III):** the AI falls into one of eight listed use-case categories (biometrics, critical infrastructure, education, employment, essential services, law enforcement, migration, justice). If neither route is triggered, the default is minimal risk.

4. The GPAI Layer — Foundation Models

General-Purpose AI (GPAI) models — the large foundation models like GPT, Claude, and Gemini that underpin many applications — are governed under a separate chapter (Chapter V) that sits alongside, not inside, the risk tier framework.

All GPAI providers face baseline obligations regardless of risk: technical documentation, training data summaries, and copyright compliance policies. The heavier obligations apply to models classified as having **systemic risk** — currently defined as models trained using more than 10^{25} floating point operations (FLOPs).

GPAI providers with systemic risk must additionally conduct adversarial testing, assess and mitigate systemic risks, report serious incidents to the AI Office, and maintain cybersecurity protections.

Critical point for practitioners: when you build an application on top of a GPAI model, the GPAI provider remains responsible for their model's obligations. You become responsible for the AI system you create on top of it. These are separate layers of accountability.

5. When Does It Apply? — The Timeline

The Act applies in phases, not all at once.

Date	What happens
2 Feb 2025	Prohibited practices (Chapter I & II) take effect — the bans are live now
2 Aug 2025	GPAI model obligations apply (Chapter V, VII, XII) — foundation model providers must comply

Date	What happens
2 Aug 2026	Full application — high-risk AI system obligations, governance, market surveillance
2 Aug 2027	Article 6(1) high-risk systems (those embedded in regulated products under Annex I) fully apply

Practical implication: as of March 2026, the prohibited practices are already in force and GPAI obligations are live. The full high-risk framework for most systems applies from August 2026 — less than six months away at time of writing.

6. What Are the Consequences?

The Act sets penalties at a scale designed to be meaningful even for the largest organisations.

Infringement type	Max fixed fine	Or % of global turnover
Prohibited practices (Article 5)	**€35,000,000**	7% (whichever is higher)
High-risk system non-compliance	**€15,000,000**	3% (whichever is higher)
GPAI provider non-compliance (Article 101)	**€15,000,000**	3% (whichever is higher)
Supplying incorrect information to authorities	**€7,500,000**	1% (whichever is higher)

Fines are set at the *higher* of the fixed amount or the percentage of global annual turnover — meaning the percentage clause is most consequential for large organisations, while the fixed amounts set a meaningful floor for everyone else.

7. The One Question to Ask First

Before any conversation about AI compliance, obligations, or risk — this is the question that determines everything else:

The diagnostic question for every AI engagement:

"Is this deliverable an AI system under the Act's definition — and if so, am I the provider or the deployer of it?"

These are three separate questions:

- Does the system perform dynamic inference — or is it a static tool or work product created with AI assistance?
- If it is an AI system, which risk tier does it fall into?
- Am I the provider (built and deployed it under my name) or the deployer (using a third-party system)?

Getting these three answers right is the foundation of every compliance decision that follows.

8. What This Means for You

The Act is real, it is in force, and its penalties are significant. But for most people reading this, the honest answer to what do I need to do is considerably less dramatic than the headlines suggest. The right response depends entirely on your role.

If you use AI tools at work

Your existing tools are almost certainly not illegal. Consumer and enterprise AI products such as Microsoft Copilot, ChatGPT, Claude, and Gemini are developed by providers who carry the compliance obligations for those platforms. Using them does not make you a provider under the Act.

What is worth checking: whether the tools your organisation has approved are enterprise-tier rather than consumer-tier. This matters less for Act compliance and more for data governance, but it is the same conversation. If your organisation has no AI tool policy at all, that gap is worth raising. The Act does not require you to stop using AI. It requires that AI affecting people's rights or safety is used responsibly.

If your organisation deploys AI

Start with the diagnostic question from Section 7. Work through your AI systems one by one: is this system performing dynamic inference, or is it a static tool or AI-assisted work product? If it is an AI system, which risk tier does it fall into? Most operational AI tools will land in minimal risk. No mandatory obligations.

Where it becomes material is if your organisation uses AI in hiring decisions, customer credit assessments, healthcare eligibility, or any of the Annex III categories. If so, you are likely deploying a high-risk system and deployer obligations apply from August 2026. The time to understand those obligations is now, not in July.

If you build or configure AI systems for others

This is where the Act's weight falls most directly. If you develop an AI system, including building on top of a foundation model, and deliver it to a client under your own name, you are a provider. Provider obligations are not light for high-risk systems: technical documentation, conformity assessment, registration, and human oversight mechanisms.

The critical first step is not drafting compliance documentation. It is correctly classifying what you are building. A misclassification, treating a high-risk system as minimal risk, is itself a compliance failure. If you regularly build AI tools for clients, getting that classification methodology right is the foundation everything else rests on.

The bottom line

The EU AI Act does not require you to stop using AI. It does not require most people to file paperwork with the EU. What it requires is that the people building and deploying AI systems take their obligations seriously, and that organisations understand which category they fall into.

The four-lens analysis that accompanies this guide examines the Act in detail across four specific audiences: pharma and HEOR, AI strategy consultants, SMEs, and the general governance landscape. If one of those lenses is yours, that is the right next read.

Want to discuss what the EU AI Act means for your organisation?

Get in touch: ryan@ryanbishop.co.uk

Source

This guide is based on the full text of Regulation (EU) 2024/1689 as published in the Official Journal of the European Union.

EU AI Act full text: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

Key articles referenced: Article 3 (definitions), Article 5 (prohibited practices), Article 6 (high-risk classification), Annex I (harmonisation legislation), Annex III (high-risk use cases), Article 50 (transparency obligations), Articles 51–56 (GPAI), Article 99 (penalties), Article 101 (GPAI fines), Article 113 (entry into force).

Want to discuss what the EU AI Act means for your organisation? Get in touch: ryan@ryanbishop.co.uk