AI LITERACY SERIES

The EU AI Act: A Four-Lens Analysis

Pharma & HEOR | AI Strategy Consultant | SME | General Governance

Ryan Bishop | March 2026 | ryanbishop.co.uk

Introduction

The EU AI Act is a single piece of legislation. But what it means in practice depends entirely on who you are and what you do. A frontier AI developer, a solo consultant, a dental practice, and a pharmaceutical company all face the same law — and arrive at entirely different answers about what it requires of them.

This analysis applies a critical thinking framework to the Act across four distinct lenses, moving beyond description to ask the harder questions: why does this matter, whether the framework works as intended, and what the verdict is for each audience.

Each lens follows the same structure: what matters for this audience, the key provisions at play, the practical implications, and a verdict. Readers who have not yet read the accompanying plain-English cheat sheet may wish to start there for context on the Act's structure and terminology.

A note on methodology: *The analysis presented here draws on the Act's full text (Regulation EU 2024/1689). Where conclusions are contested or uncertain, that uncertainty is made explicit. The purpose is critical assessment, not advocacy.*

LENS 1

Pharma & HEOR

Where does AI in the health technology lifecycle actually land?

What Matters

The Act's implications differ significantly depending on where in the product lifecycle AI is being used. Three distinct positions need to be considered separately.

Position 1 — AI in medical devices and Software as a Medical Device (SaMD)

This is the most clearly regulated position. Article 6(1) combined with Annex I creates a direct high-risk pathway for AI that is a safety component of a product subject to Union harmonisation legislation. The Medical Devices Regulation (MDR) and In Vitro Diagnostic Regulation (IVDR) are both listed in Annex I. AI-powered diagnostic tools, SaMD, and AI used as a safety component of medical devices are **high-risk by definition** — not by use case assessment, but by structural classification. Full Chapter III obligations apply.

Position 2 — AI in drug discovery and development

More nuanced. A target identification model or molecular design tool is not itself a medical device. As a standalone analytical tool it likely sits outside both Annex I and Annex III, placing it in minimal risk. However, the boundary shifts if AI output feeds directly into a regulated process — clinical trial design, regulatory submission data, safety signal detection. The key question is whether the AI is influencing a regulated decision, and how traceable that influence is.

Position 3 — HEOR and HTA submission preparation

This is where the most analytically significant finding emerges — and where the Act's framework is most revealing about its own limitations.

Key Provisions

HTA submission preparation — economic modelling, systematic literature reviews, network meta-analyses, value dossier writing — sits outside the Act's high-risk classification. Working through both routes to high-risk status makes this clear.

The Annex I route does not apply. These are analytical instruments used by private organisations, not safety components of regulated medical devices.

The Annex III route does not apply. The closest healthcare trigger requires AI used by *public authorities* to evaluate *individual* eligibility for healthcare services. HTA preparation fails on both counts: the deployer is a private entity, and the analysis operates at population level with no individual natural person being assessed.

The Act's classification answer is therefore: **minimal risk, no mandatory obligations.**

Practical Implications

That classification is technically correct under the Act as written. But three specific use cases expose why it is deeply uncomfortable:

Digital twins and synthetic trial arms for rare disease

If used in a clinical trial context, synthetic patient data potentially enters via Annex I through the Clinical Trials Regulation. Even outside that context, a synthetic trial arm built to represent a rare disease population of 200 patients across Europe is functionally operating on a near-identifiable group. The abstraction between population-level analysis and individual natural persons collapses at that scale.

Black box algorithms in cost-effectiveness models

The ICER is the central output of an HTA submission. An opaque algorithm producing that output would fail the Act's own transparency and explainability principles for high-risk systems — but only if it were classified as high-risk in the first place. The current framework does not reach it.

AI hallucination in systematic literature reviews

If the evidence base underpinning a submission is corrupted by hallucinated references or misrepresented findings, the entire value dossier may rest on false foundations. HTA bodies make reimbursement decisions on that evidence. Patients access or are denied treatment on the basis of those decisions.

These examples expose a structural consequence of the Act's architecture, not a drafting oversight. The Act is built around protecting natural persons from consequential AI decisions made *about* them. In HTA preparation, the natural person never appears in the AI interaction — they only appear downstream, in the HTA body's deliberation.

The enforcement gap

HTA bodies are developing AI guidance — NICE, EMA, and EUnetHTA have all produced or are producing reflection papers on AI use in evidence generation. But HTA body guidance regulates the *submission*. The Act, if it applied, would regulate the *tool*. These are different regulatory layers, operating on different actors. HTA bodies cover the output. The Act would cover the process.

The only instrument with real legal teeth — fines up to 7% of global turnover — does not reach this part of the AI lifecycle. Existing regulatory guidance has principles but no enforcement mechanism.

VERDICT

The Act's individual-centric rights framework creates a structural blind spot around population-level AI use in healthcare evidence generation — precisely where AI failure carries the greatest systemic risk to patient access, and where existing regulatory guidance has principles but no teeth.

This is not a criticism of the Act's intent. It is a consequence of building a rights framework around individual persons in a domain where some of the most consequential AI sits one step removed from any identifiable individual. Addressing this gap will require either an amendment to the Act's scope, or HTA bodies acquiring enforceable AI governance powers of their own.

LENS 2

AI Strategy Consultant

Provider, deployer, or neither — and why getting this wrong matters.

What Matters

For an AI strategy consultant, the Act presents an unusual challenge: you are simultaneously the analyst advising clients on compliance, and a potential subject of that compliance yourself. The two roles carry different obligations — and the boundary between them is less obvious than it first appears.

The central question is not whether you use AI in your work. Almost every knowledge professional does. The question is what you do with it, and what you deliver as a result.

Key Provisions — The Provider / Deployer Distinction

The Act's obligations are asymmetric by design. Providers carry the heaviest burden — technical documentation, conformity assessment, registration. Deployers carry lighter but real obligations — human oversight, transparency, logging, incident reporting for high-risk systems.

Four practical scenarios illustrate how the classification plays out for a consultant:

Scenario 1 — A static flashcard or reference tool

A tool that displays pre-authored content without performing dynamic inference is **not an AI system** under the Act. The Act defines an AI system as one that infers outputs from inputs. No inference, no obligation. The fact that AI was used to help create the content is irrelevant — the Act regulates AI systems in deployment, not work products created with AI assistance.

Scenario 2 — A semantic retrieval or recommendation tool

A system that takes a document brief as input and recommends relevant references using embedding-based similarity is performing inference — and is therefore an AI system under the Act. However, working through both Annex I and Annex III routes, it almost certainly lands in **minimal risk**. No mandatory compliance obligations. But if that tool is built and delivered to a client under the consultancy's name, the consultant is a **provider** of that AI system — with all that entails, should the risk classification ever change.

Scenario 3 — Building a client website with embedded third-party tools

A website containing Google Maps and a booking system is not an AI system. The consultant building it is operating as a developer and service provider — not an AI provider or deployer under the Act. Responsibility for compliance of the embedded third-party tools rests with their respective developers.

Scenario 4 — Creating a document or presentation using GenAI

The deliverable is a document. A document is not an AI system. No Act obligations apply to the consultant. The caveat is context-specific: if the document is a pharmaceutical HTA submission, the consultant's pharma client may face growing expectations from HTA bodies to declare AI involvement in submission preparation — not because the Act requires it, but because HTA governance frameworks are developing independently.

Practical Implications

The pattern across all four scenarios is consistent, and it produces a principle worth committing to:

The diagnostic question for every client engagement

"Is the 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. Getting them right before any client engagement begins is the foundation of every compliance and contractual decision that follows. The consultant role and the provider role carry different legal relationships — and a contract that does not reflect which one applies is a liability.

There is a broader point worth making about professional standards. As a consultant you are rarely a provider under the Act. But the Act's influence on your practice is real, because it shapes the compliance questions your clients will bring to you — and the expectations they will carry about AI transparency in the work you deliver on their behalf.

This matters particularly in pharmaceutical and healthcare contexts. A pharma client receiving an AI-assisted deliverable may face HTA body expectations around declaring AI involvement, even where no legal obligation currently exists. Being ahead of that conversation is a differentiator, not a burden.

VERDICT

As a consultant you are rarely a provider under the Act — but your clients' obligations, and their relationships with their own regulators, mean AI transparency in your work is a professional standard question even where it is not a legal one.

The consultant who can answer the classification question clearly, advise clients on where their obligations begin, and build that thinking into contract structures and ways of working is offering something the Act itself does not: navigability.

LENS 3

Small Medium Enterprise (SME)

The Act reduces the cost of compliance. It does not reduce the complexity.

What Matters

The Act is unusual in that it names SMEs explicitly. Article 62 creates specific concessions: priority access to regulatory sandboxes, reduced conformity assessment fees, specific awareness-raising and training activities. Article 63 allows microenterprises to comply with quality management requirements in a simplified manner.

On the surface, this reads as a considered accommodation of smaller operators. Examined more carefully, it exposes a significant gap between what the Act offers SMEs and what SMEs actually need.

Key Provisions

The Act's SME concessions cluster around three areas: access, cost, and simplification. All three share a critical assumption — that the SME has already correctly identified that it has a high-risk AI system requiring conformity assessment. That assumption does not hold in practice.

Consider a dental practice. Dental AI tools — diagnostic imaging analysis, caries detection, bone loss assessment from radiographs, treatment planning support — are operating on individual patients in a clinical setting, influencing clinical decisions. Under Article 6(1) combined with Annex I, AI used as a safety component of a medical device subject to the MDR is **high-risk by definition**.

As a deployer of that high-risk system, the dental practice carries real obligations:

- Conduct a fundamental rights impact assessment before deployment
- Ensure human oversight mechanisms are documented and in place
- Maintain logs of AI system outputs for a minimum of six months
- Report serious incidents to market surveillance authorities
- Verify the system carries an EU declaration of conformity before use
- Inform patients that AI is being used in their care

The same analysis applies to physiotherapists using AI movement analysis tools, GPs using AI diagnostic support, and optometrists using AI retinal screening systems. The dental practice is not an edge case — it is representative of an entire category of healthcare SME that the Act reaches directly.

Practical Implications

Here is the structural problem: the practice owner deploying an AI imaging system almost certainly does not know any of this.

The Act's concessions — sandbox priority access, reduced conformity assessment fees — are entirely irrelevant to that practice owner. They have not reached the question of fees. They have not correctly identified that they have a compliance obligation at all. The concessions address affordability. They do not address the prior question of whether the operator understands they are subject to the Act in the first place.

The navigability gap

The classification decision itself — working through Article 6, Annex I, Annex III, and the Article 6(3) carve-outs — requires legal and technical expertise that a small operator is least likely to have. A large organisation misclassifying an AI system as minimal risk when it is high-risk has infrastructure to catch the error. An SME making the same misclassification has no safety net.

The penalties do not scale with company size in a way that changes this dynamic. A fine of up to 3% of global annual turnover — or €15M, whichever is higher — is existential for an SME in a way it is not for a large enterprise.

There is also a temporal problem. The Act's concessions assume an SME that is actively engaging with compliance questions. The dental practice deploying AI imaging in August 2026 may not have given the Act a moment's thought. Market surveillance authority activity will initially focus on larger, more visible operators. But the structural exposure is real, and the timeline is not distant.

VERDICT

The Act's SME concessions address affordability. They do not address navigability. For an SME, getting the classification wrong in the first place is the existential risk — and the Act provides no meaningful help with that decision.

The practical implication for any adviser working with SMEs is clear: the most valuable thing you can do is not help them prepare compliance documentation. It is help them answer the classification question correctly before any other decision is made.

LENS 4

General Governance

Good law. Untested enforcement. Racing against its own obsolescence.

What Matters

The EU AI Act presents itself as three things simultaneously: a safety framework protecting people from harmful AI, an innovation framework enabling responsible development, and a values framework embedding human-centric principles into AI governance. This is an ambitious combination.

The Act's own recitals are clear: safety is the stated primary objective. The question is whether the Act's architecture actually reflects that priority.

Key Provisions — Testing the Safety-First Claim

Three features of the Act's structure warrant scrutiny.

The prohibited practices list is the easy consensus

Article 5 bans social scoring, subliminal manipulation, and most real-time biometric surveillance. These prohibitions are important. They are also, largely, practices that were already widely condemned before the Act existed. The Act codified consensus rather than staking out contested ground. The harder cases — the structural blind spots identified in Lens 1, the navigability failures identified in Lens 3 — are not addressed by the prohibited list.

High-risk conformity assessment is predominantly self-declared

For the majority of Annex III high-risk AI systems (points 2 to 8), the conformity assessment procedure is internal control — no third-party body involved. A provider self-certifies compliance and places the system on the market. Compare this to the pharmaceutical regulatory model, where no drug reaches patients without independent review. The safety parallel the Act implicitly draws does not hold structurally.

The GPAI chapter is softer — but not toothless

A first reading of the GPAI chapter might suggest that the most powerful AI systems face the lightest touch. That reading requires correction. Providers of GPAI models with systemic risk face mandatory adversarial testing, systemic risk assessment and mitigation, incident reporting within two weeks, and cybersecurity obligations. The AI Office holds direct enforcement jurisdiction — not delegated to 27 member states. Fines up to €15M or 3% of global turnover apply under Article 101.

The more precise observation is not that large AI providers face softer obligations, but that their obligations are governed through a different and less mature enforcement architecture. The AI Office is new. Its capacity to monitor frontier model providers at the required scale remains to be demonstrated.

Practical Implications — The Enforcement Question

The GDPR comparison is instructive. The regulation came into force in May 2018. The first major fine did not land until October 2020 — two and a half years of theoretical obligations before enforcement reality caught up with legislative intent. And GDPR was regulating data handling, a relatively static concept that organisations already understood.

The AI Act faces a harder version of the same problem. It is regulating a technology that is actively evolving faster than the regulatory cycle. The compute threshold for GPAI systemic risk — 10^{25} floating point

operations — is already looking like a blunt instrument as algorithmic efficiency improves. The Act acknowledges this, empowering the Commission to amend thresholds via delegated acts. But that is a reactive mechanism chasing a proactive challenge.

The obsolescence question

By the time enforcement capacity exists at the scale required, the technology the Act was written to govern may have changed sufficiently to make parts of it obsolete. This is not a criticism unique to the EU AI Act — it is a structural challenge for any attempt to regulate technology through legislation.

The Act's delegated acts provisions and the Commission's power to amend Annex III are designed to address this. Whether they are agile enough in practice remains the open question.

None of this diminishes the Act's significance. It is the world's first comprehensive AI governance framework with genuine legal teeth. The prohibited practices are in force now. The GPAI obligations are live. The high-risk framework applies from August 2026. This is real regulation, not a voluntary code.

But the honest assessment of a genuinely serious piece of legislation is not uncritical endorsement. The Act is right to exist. Its architecture makes reasonable compromises given the political and technical environment in which it was drafted. Whether those compromises hold as the technology evolves is the question that will define its legacy.

VERDICT

Good law. Untested enforcement. Racing against its own obsolescence.

The EU AI Act is a landmark piece of legislation — the right idea, seriously designed, with genuine enforcement mechanisms. It faces a dual challenge no regulation has yet successfully solved: building enforcement capacity faster than public harm can occur, while remaining relevant to a technology that evolves faster than the legislative cycle. Judging it on its intent alone would be too generous. Dismissing it for its limitations would be too cynical. The appropriate assessment is to take it seriously, comply with what it requires, and remain alert to where it does not yet reach.

Source

This analysis is based on the full text of Regulation (EU) 2024/1689 as published in the Official Journal of the European Union, cross-referenced against the Cottrell critical thinking framework.

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

Key articles and annexes referenced: Article 3 (definitions), Article 5 (prohibited practices), Article 6 (high-risk classification), Article 7 (amendments to Annex III), Articles 8-15 (high-risk requirements), Article 26 (deployer obligations), Article 50 (transparency obligations), Articles 51-56 (GPAI models), Article 62 (SME measures), Article 63 (microenterprise derogations), Article 99 (penalties), Article 101 (GPAI fines), Article 113 (entry into force), Annex I (harmonisation legislation), Annex III (high-risk use cases).

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