



AI Literacy Lab



Bykov-Brett
Enterprises

EU AI ACT / OBLIGATIONS BRIEF

In Scope

What the EU AI Act now requires of your business

Obligations, deadlines and penalties, explained without the legalese.

Turn Intent Into Impact.



The date everyone is braced for has moved

In July 2026 the EU changed what happens on 2 August. If you are preparing for high risk obligations in August, you are preparing for the wrong thing.

STILL LANDS 2 AUGUST 2026

Transparency

- Tell people when they are talking to a machine.
- Mark AI generated output so a machine can read the mark.
- Disclose deepfakes.
- The Commission can fine general purpose AI model providers for the first time.

Transparency is what bites for most businesses, but it is not all that lands. The conformity assessment and registration machinery also applies from this date, even though the duties it attaches to do not apply until December 2027.

WHAT MOVED, AND BY HOW LONG

High risk duties

High risk systems on the Annex III list, such as hiring, credit, education

2 August 2026 becomes 2 December 2027

High risk AI built into regulated products on Annex I

2 August 2027 becomes 2 August 2028

Regulation (EU) 2026/1744, the Digital Omnibus on AI, Article 1(40). Signed 8 July 2026, published in the Official Journal 24 July 2026, in force 27 July 2026.

You do not have to sell in the EU to be caught by this

"providers and deployers of AI systems that have their place of establishment or are located in a third country, where the output produced by the AI system is used in the Union"

Article 2(1)(c). The Act says its purpose here is "to prevent the circumvention of this Regulation".

UK SAAS, EU CUSTOMER

You sell an AI meeting summariser from Manchester. One customer is in Rotterdam. On first making it available, you have placed an AI system on the Union market. A free beta counts the same.

Article 2(1)(a)

UK RECRUITER, DUBLIN ROLE

You rank CVs on UK servers to shortlist for a Dublin employer. The software never leaves the UK. The shortlist is acted on in Dublin, so the output is used in the Union.

Article 2(1)(c)

UK PARENT, EU SUBSIDIARY

Your UK credit model scores applicants. Your Frankfurt subsidiary lends on those scores. Two hooks fire independently, and the subsidiary cannot argue its way out.

Article 2(1)(b) and (c)

GET THIS RIGHT The test is where the output is used, not where the person lives. A UK agency screening Irish candidates for a UK role is not obviously in scope.

Four levels of risk, and only two of them cost you anything

PROHIBITED Article 5	Uses banned outright. No safeguard makes them lawful.	Stop. There is no compliance route.
HIGH RISK Article 6, Annexes I and III	Systems that decide or materially shape something consequential about a person: a job, a loan, a grade, a benefit.	Risk management, data governance, documentation, logging, human oversight, conformity assessment, registration.
LIMITED RISK Article 50	Not dangerous, but people must not be deceived. Chatbots, deepfakes, synthetic media.	Disclose. Label AI output so a machine can read the mark.
MINIMAL RISK No obligations imposed	Everything else, and it is the vast majority. Spam filters, recommenders, forecasting.	Nothing mandatory, with one exception below.

THE OBLIGATION THAT IGNORES THE TIERS

Article 4 applies to every provider and deployer of any AI system, at any tier, of any size, in any sector. A company whose entire AI footprint is minimal risk still has a live legal duty.

A note on language: the four levels are the Commission's way of explaining the Act, not words used in it. The operative text creates prohibitions, high risk duties and transparency duties. Useful as a map, not as a legal category.

Some uses are illegal, whatever safeguards you put around them.

These have been prohibited since 2 February 2025 and carry the heaviest penalty tier. There is no compliance route, no safeguard and no consent that makes them lawful.

a Subliminal or manipulative techniques that distort behaviour and cause significant harm

b Exploiting vulnerabilities of age, disability, or social and economic situation

c Social scoring: rating people on behaviour or traits, then penalising them in unrelated contexts

d Predicting criminality from profiling or personality traits alone

e Untargeted scraping of facial images to build recognition databases

f **Emotion inference in the workplace or in education, with medical and safety exceptions only**

g Biometric categorisation to infer race, politics, union membership, religion, sex life or sexual orientation

h Real time remote biometric identification in public spaces for law enforcement, with narrow exceptions

POINT (F) IS THE ONE THAT CATCHES ORDINARY EMPLOYERS

Sentiment analysis on staff surveys, and emotion scoring on call centre recordings, fall under this ban. There is no version of them you can run safely, however careful the process around them.

NEW, FROM 2 DECEMBER 2026

Two further prohibitions are now law: AI generating realistic intimate material of an identifiable person without consent, and AI generating child sexual abuse material. Both catch use, not only supply, so deployers are caught too.

The eight high risk areas, and what each looks like in practice

1	Biometrics	Face recognition identifying people from CCTV at scale. Verification at a turnstile is excluded.
2	Critical infrastructure	AI managing traffic light timing, or water and electricity supply.
3	Education and vocational training	Exam proctoring software monitoring a candidate by webcam.
4	Employment and worker management	CV screening and candidate ranking.
5	Essential public and private services	Credit scoring for a loan application. Fraud detection is excluded.
6	Law enforcement	Predictive reoffending risk scoring.
7	Migration, asylum and border control	Automated triage of visa applications.
8	Justice and democratic processes	AI drafting judicial reasoning, or influencing voting behaviour.

THE THREE HIGHLIGHTED ROWS ARE THE ONES A NORMAL COMPANY TOUCHES

Areas 1, 6 and 7 apply only in so far as the use is permitted under other law. Areas 3, 4 and 5 carry no such qualifier. Education, employment and credit are where an ordinary business finds itself unexpectedly in scope.

"an AI system referred to in Annex III shall always be considered to be high-risk where the AI system performs profiling of natural persons"

Article 6(3). There is an exemption route out of high risk, and profiling closes it absolutely.

Eleven duties you can run from a register and a named owner

These bind you as a user of AI you did not build. They apply from 2 December 2027 for Annex III systems, which is not as long as it sounds if you have to change how decisions get made.

1 Use it as the instructions say
Article 26(1)

2 Name real people to oversee it, with the competence, training and the authority to overrule it
Article 26(2)

3 Check the input data, where you control it
Article 26(4)

4 Monitor it, and suspend use if it starts presenting a risk
Article 26(5)

5 Keep the logs for at least six months
Article 26(6)

6 Tell workers and their representatives before you switch it on
Article 26(7)

7 Tell the people it makes, or helps make, decisions about
Article 26(11)

8 Cooperate with the authorities
Article 26(12)

A MYTH WORTH KILLING

You will read that anyone deploying high risk AI must run a fundamental rights impact assessment. That is wrong. It binds public bodies, private organisations providing public services, and deployers doing credit scoring or life and health insurance pricing. A private company using AI to hire owes every duty above and no FRIA.

Article 27(1)

Buying it does not make you the provider. Branding it does.

Three things turn a user of AI into its legal provider. None of them require you to build anything.

YOU PUT YOUR NAME ON IT

You white label a high risk system and ship it as your own. Your vendor contract can settle who pays. It cannot settle who the regulator holds responsible.

Article 25(1)(a)

YOU REBUILT PART OF IT

A change not foreseen in the original conformity assessment. Fine tuning on your own data, rewiring decision logic, bolting on a scoring layer.

Article 25(1)(b)

YOU USED IT FOR SOMETHING ELSE

You narrow a general purpose tool to a high risk purpose. This is the one that catches ordinary companies.

Article 25(1)(c)

"the provider that initially placed the AI system on the market or put it into service shall no longer be considered to be a provider of that specific AI system"

Article 25(2). The obligation is not shared. It transfers.

Article 26 is eleven duties you can run from a register and a named owner. Article 16 is a quality management system, a conformity assessment and a CE mark. Article 25 is the door between them.

Binding for eighteen months, and the omnibus left it alone

BINDING SINCE

2 February 2025

SOFTENED

27 July 2026

ENFORCED FROM

2 August 2026

UNTIL 26 JULY 2026

"shall take measures to ensure, to their best extent, a sufficient level of AI literacy of their staff"

An obligation of result. You had to reach a level.

FROM 27 JULY 2026

"shall take measures to support the development of AI literacy of their staff ... This obligation does not require providers or deployers to guarantee any specific level of AI literacy of any individual."

An obligation of effort. You have to try, and record that you tried.

"AI literacy should be a strategic priority, regardless of regulatory obligations and potential sanctions."

Regulation (EU) 2026/1744, recital 8. Brussels relaxed the rule and said in the same breath that you should do it anyway.

Worth knowing: Article 4 carries no EU level fine. It sits outside the Article 99(4) tier. The provisions with teeth are Article 26(2), which requires named human overseers with the competence, training and authority to intervene, and Article 50.

From 2 August 2026, you have to say when it is AI

This part arrives on 2 August 2026. It applies whatever tier your system sits in, and it stacks on top of high risk duties rather than replacing them.

AI THE BASIC MARK Whenever AI had a hand in making the content, or as a custom label such as "voice generated with".	AI GENERATED FULLY AI GENERATED The whole piece is AI made, with no human input beyond the prompt.	AI MODIFIED PARTIALLY AI MODIFIED Real human made content, altered by AI, on a matter of public interest.
Tell people when they are interacting with an AI system, unless it is obvious Article 50(1)	Mark it so a machine can read that the output is synthetic Article 50(2)	Disclose deepfakes and AI text published to inform the public on matters of public interest Article 50(4)
DO NOT OVERREAD THE GRACE PERIOD You may read that transparency slipped to December 2026. It did not. Only the machine readable marking duty gets extra time, only for providers, and only for systems already on the market before 2 August. Telling people they are talking to AI, and disclosing deepfakes, start on 2 August 2026 with no grace at all.		

The marks above are shown to illustrate how disclosure is being standardised in practice. The legal duty is to disclose and to mark machine readably. The Act does not prescribe a specific icon.

Calling an API does not make you a model provider

You are the provider of the system you built, not of the model underneath it. The duties split cleanly, and most companies carry far less than they fear.

SITS WITH THE MODEL PROVIDER

OpenAI, Anthropic, Google, Mistral

Technical documentation of training and evaluation. An information pack for everyone integrating the model. A copyright compliance policy. A public summary of the training content. Extra duties again if the model carries systemic risk.

Article 53, in force since 2 August 2025

SITS WITH YOU

The company building the product

The Article 5 prohibitions, which are already live. Transparency under Article 50 from 2 August 2026. AI literacy under Article 4. And the full high risk regime, but only if your use case is on the Annex III list.

Your nearest exposure is Article 50

WHAT IF WE FINE TUNE IT

You only become the provider of the model itself if the compute used for your modification exceeds one third of the compute used to train the original. The Commission's own guidance adds that currently few modifications meet that criterion.

Ordinary fine tuning on a frontier model is nowhere near that. If you do cross it, you also have to appoint an authorised representative established in the EU.

Commission Guidelines on the scope of obligations for providers of general purpose AI models, C(2025) 7719 final, 19 November 2025.

The number everyone quotes is the wrong one for most companies

Breach	Cap	Turnover	Which applies	Article
Prohibited AI practices	EUR 35m	7%	the HIGHER	99(3)
Operator obligations, including providers (Art 16), deployers (Art 26) and transparency (Art 50)	EUR 15m	3%	the HIGHER	99(4)
Supplying incorrect or misleading information to authorities	EUR 7.5m	1%	the HIGHER	99(5)

IF YOU ARE AN SME OR A START UP, READ THIS TWICE

The test reverses. Article 99(6) makes your fine the LOWER of the two figures, not the higher. The Digital Omnibus extends the same treatment to small mid-caps.

Almost every summary prints "whichever is higher" and stops there. For a small company that overstates the exposure enormously. The percentage also applies only where the offender is an undertaking.

The honest picture: no monetary fine has been publicly issued under the AI Act. As of 17 June 2026 only 9 of 27 Member States had designated both a market surveillance and a notifying authority, against a deadline of 2 August 2025. The penalties have been legally available since 2 August 2025. The gap is capacity, not authority.

Turnover means total worldwide annual turnover for the preceding financial year, group wide, not EU revenue and not the AI product's revenue. Figures from Regulation (EU) 2024/1689 Article 99, as amended by Regulation (EU) 2026/1744.

EU AI Act – The whole runway, in one place

● 2 Feb 2025	PAST	Prohibited practices became illegal. The AI literacy duty started.
● 2 Aug 2025	PAST	Obligations for general purpose AI models. Penalties became legally available.
● 27 Jul 2026	IN FORCE	The Digital Omnibus entered into force and moved the high risk deadlines.
● 2 Aug 2026	TRANSPARENCY LANDS	Transparency duties. The conformity and registration machinery. The Commission can fine model providers.
● 2 Dec 2026	NEW PROHIBITIONS	Two new prohibitions. Marking duty for synthetic content systems already on the market.
● 2 Aug 2027		Legacy general purpose models placed on the market before August 2025 must comply.
● 2 Dec 2027	THE BIG ONE	High risk duties for Annex III systems: hiring, credit, education, essential services.
● 2 Aug 2028		High risk duties for AI embedded in regulated products under Annex I.
● 2 Aug 2030		Public authority legacy systems must comply.

Dates for the amended provisions come from Regulation (EU) 2026/1744 Article 1(40). There is no consolidated version of the AI Act yet, so a consolidated Article 113 cannot be cited.

Ten questions. Every one you cannot answer is a piece of work.

1 Can you list every AI system in use across the business, including the ones people adopted themselves?
Your register

2 For each one, do you know whether you are the provider or the deployer?
Article 3

3 Has anyone checked whether any of them sits on the Annex III list?
Article 6

4 Is anything inferring emotion from staff or students?
Article 5(1)(f)

5 Do people affected by an AI assisted decision get told?
Articles 26(11) and 50

6 Where a person oversees an AI system, do they have the authority to overrule it?
Article 26(2)

7 Can you show what you have done to build AI literacy, and when you did it?
Article 4

8 Does anything you publish that was AI generated say so?
Article 50(4)

9 Would you know if someone repurposed an existing tool for a new job without telling you?
Article 25(1)(c)

10 If a regulator asked tomorrow, which document would you hand them?
All of the above

Few people answer all ten on the first pass. The score does not matter. The first three questions are a morning's work, and almost nothing else can start until they are done.

WHERE TO START

Start with the duty that has been law the longest

AI LITERACY THAT HOLDS UP TO AUDIT

Short, accessible e-learning built for exactly the duty on page nine. Records what was covered, for whom, and when.

ailiteracylab.ai

A CONVERSATION ABOUT YOUR REGISTER

If the first three questions on the previous page are the problem, that is where we start.

bykovbrett.net

IF YOU CHECK THIS AGAINST THE OFFICIAL SOURCE

You will see wording that contradicts this brief. There is no consolidated version of the AI Act yet, and the Commission's own pages still carry the superseded text. Every provision here is cited to Regulation (EU) 2026/1744 directly, so you can check the amending law itself.



ailiteracylab.ai/resources

Turn Intent Into Impact.



Researched and drafted with AI, under human editorial responsibility.