



Prohibited AI practices

All ten prohibitions, each with the threshold that engages it, the carve-outs written into the paragraph, and the lawful practice it is most often confused with.

10

prohibited practices

8

in force today

2

added by the Omnibus

The practices that have no compliance route

A prohibition is the only finding in the Act that cannot be planned around: there is no conformity assessment, no documentation package and no mitigation that makes it lawful. The practice stops. Eight have bound since 2 Feb 2025.

Where a prohibition sits on the risk ladder

Prohibited No conformity route. The practice stops or changes.
High-risk Permitted, with the full Chapter III obligations.
Transparency Permitted, with a disclosure duty under Article 50.
Minimal No system-specific requirements beyond general law.

The ladder is read from the top down, and a finding stops the descent. Everything below a prohibition is assessment work that a prohibited system never reaches.

Which of the ten bind today



■ In force since 2 February 2025 · 80% ■ Added by the Omnibus, from 2 December 2026 · 20%

Any material listing eight prohibitions is describing the superseded text.

5(1)(a) — Subliminal, manipulative or deceptive techniques

AI systems that deploy subliminal, purposefully manipulative or deceptive techniques to materially distort a person's behavior are prohibited, but only where that distortion causes or is reasonably likely to cause significant harm.

Threshold. Two conditions, both required: the technique must materially distort behavior by appreciably impairing the ability to make an informed decision, and that distortion must cause or be reasonably likely to cause significant harm to that person, another person or a group.

Confused with. Persuasive design, dark patterns and recommender optimization — Generally outside Article 5 unless the significant-harm threshold is crossed, but squarely within consumer protection law, the Digital Services Act and the GDPR

Carve-outs: Lawful persuasion that does not impair informed decision-making; Ordinary advertising and personalisation that falls short of material distortion; Approved therapeutic uses carried out with informed consent.

Article 5(1)(a)

5(1)(b) — Exploiting vulnerabilities of age, disability or circumstance

AI systems that exploit vulnerabilities arising from a person's age, disability, or specific social or economic situation in order to materially distort their behavior are prohibited where that causes or is reasonably likely to cause significant harm.

Threshold. The system must exploit a vulnerability of one of the named kinds, materially distort behavior, and cause or be reasonably likely to cause significant harm. Note that economic hardship is expressly a listed vulnerability.

Confused with. Age-based or income-based audience targeting in advertising — Not prohibited by itself. The ban requires exploitation of the vulnerability plus behavioral distortion plus significant harm.

Carve-outs: Accessibility and assistive technology designed to help the same groups; Targeting that does not exploit the vulnerability to distort behaviour harmfully.

Article 5(1)(b)

5(1)(ba) — Generating non-consensual intimate imagery

AI systems that generate or manipulate realistic depictions of the intimate parts, or sexually explicit conduct, of an identifiable person without that person's explicit consent are prohibited from 2 December 2026.

Threshold. The depiction must be realistic and the person identifiable, and consent must be absent. Consent has to be freely given, specific, informed, unambiguous and explicit: the GDPR standard, not a buried terms-of-service acceptance.

Confused with. General-purpose image and video generators — Not prohibited as such, but providers are liable where this misuse is a reasonably foreseeable outcome, even if unintended. The defence is demonstrable technical safeguards, so a general-purpose generator without guardrails is exposed.

Not yet in force. Applies from 2 Dec 2026.

Carve-outs: Depictions of people who are not identifiable; Use with the subject's freely given, specific, informed, unambiguous and explicit consent; Systems with reasonable and adequate technical safeguards against the outcome: the new Article 5(1a) and (1b) safeguards defence.

Article 5(1)(ba)

5(1)(bb) — Generating child sexual abuse material

AI systems that generate or manipulate child sexual abuse material within the meaning of Directive 2011/93/EU are prohibited from 2 December 2026, with only a narrow carve-out for legitimate purposes authorised under national law.

Threshold. No harm threshold and no balancing test. The material need not depict a real child: synthetic generation is caught, which closes the gap that had let purely synthetic output be argued outside existing criminal provisions.

Confused with. Safety classifiers trained to detect this material — Outside the ban: the prohibition targets generation and manipulation, not detection. Red-teaming a model to prove its safeguards hold falls within the legitimate-purpose carve-out.

Not yet in force. Applies from 2 Dec 2026.

Carve-outs: Purposes justified under national law, such as criminal proceedings; Red-teaming and safety evaluation carried out under those legitimate-purpose conditions.

Article 5(1)(bb)

5(1)(c) — Social scoring leading to unrelated detrimental treatment

AI systems that evaluate or classify people based on their social behavior or personal characteristics are prohibited where the resulting social score leads to detrimental or unfavorable treatment that is either unrelated to the context in which the data was generated, or unjustified and disproportionate to the behavior.

Threshold. The detrimental treatment must be either unrelated to the original context of the data, or unjustified or disproportionate to the social behavior. A score used for a purpose properly connected to the data, proportionately, is not caught.

Confused with. Credit scoring and insurance risk pricing — High-risk under Annex III, not prohibited. Creditworthiness assessment and life and health insurance pricing are permitted subject to the full Chapter III regime. The ban targets scores that travel across unrelated contexts.

Carve-outs: Lawful evaluation practices where the assessment is contextually related and proportionate; Sector-specific scoring that Annex III instead classifies as high-risk.

Article 5(1)(c)

5(1)(d) — Predicting individual criminality from profiling alone

AI systems that assess or predict the risk of a person committing a criminal offence based solely on profiling or on assessing their personality traits and characteristics are prohibited.

Threshold. The word doing the work is “solely”. The ban does not extend to AI that supports a human assessment already grounded in objective and verifiable facts directly linked to a criminal activity.

Confused with. Fraud detection and anti-money-laundering systems — Not caught, because these work from transactional facts rather than personality profiling. Law-enforcement risk assessment tools are separately listed as high-risk in Annex III.

Carve-outs: AI supporting a human assessment based on objective, verifiable facts directly linked to criminal activity; Place-based and offence-based analysis that does not predict individual criminality.

Article 5(1)(d)

5(1)(e) — Untargeted scraping of facial images

Creating or expanding facial recognition databases through the untargeted scraping of facial images from the internet or from CCTV footage is prohibited outright.

Threshold. No harm threshold and no balancing test. This is among the flattest prohibitions in the Article: if the scraping is untargeted and it builds or expands a facial recognition database, it is banned.

Confused with. Training a model on a licensed or consented face dataset — Outside this ban, though still subject to the GDPR and, where the system is a biometric identification system, to the Annex III high-risk regime.

Carve-outs: None stated in the paragraph itself.

Article 5(1)(e)

5(1)(f) — Emotion inference in workplaces and education

AI systems that infer the emotions of a person in the workplace or in an education institution are prohibited, except where the system is intended to be put in place for medical or safety reasons.

Threshold. Setting rather than harm is the trigger. The same emotion recognition system may be lawful in a retail context and prohibited in an office or a classroom.

Confused with. Emotion recognition in other settings — Permitted, but it triggers the Article 50(3) duty to notify the people exposed to it, and it is high-risk where Annex III applies.

Carve-outs: Systems intended for medical reasons; Systems intended for safety reasons: for example fatigue detection where alertness is a safety-critical factor.

Article 5(1)(f)

5(1)(g) — Biometric categorisation to infer sensitive characteristics

Biometric categorisation systems that categorise people in order to deduce or infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation are prohibited.

Threshold. The prohibited element is inference of one of the listed sensitive characteristics from biometric data. Biometric categorisation on other attributes is not caught by this paragraph.

Confused with. Demographic estimation for analytics — Outside this specific ban where it does not infer a listed sensitive characteristic, but still a biometric categorisation system for Annex III and Article 50(3) purposes.

Carve-outs: Labelling and filtering of lawfully acquired biometric datasets, including in the field of law enforcement; Categorisation of biometric data by law enforcement within the limits the Regulation sets.

Article 5(1)(g)

5(1)(h) — Real-time remote biometric identification in public

Real-time remote biometric identification in publicly accessible spaces for law enforcement purposes is prohibited, subject to three narrowly drawn exceptions that themselves require prior authorisation and a fundamental rights impact assessment.

Threshold. All of real-time, remote, publicly accessible space, and law enforcement purpose must be present. Post-event identification and non-law-enforcement uses fall outside this paragraph, though they may be high-risk under Annex III.

Confused with. Biometric access control at your own premises — Not a law enforcement use in a publicly accessible space, so outside this ban, but it is a biometric system with GDPR and, potentially, Annex III high-risk consequences.

Carve-outs: Targeted search for specific victims of abduction, trafficking or sexual exploitation, or for missing persons; Prevention of a specific, substantial and imminent threat to life or physical safety, or a genuine and present or foreseeable threat of a terrorist attack; Localisation or identification of a person suspected of an offence listed in Annex II and punishable by a custodial sentence of at least four years.

Article 5(1)(h)

Screen for these before anything else

HOLISTIC AI

The enterprise AI governance platform

Holistic AI governs AI across the lifecycle: an inventory of every system, the classification behind each one, the evidence a supervisory authority will ask for, and continuous monitoring once it is live.

EU AI Act readiness assessment →

holisticai.com/eu-ai-act-readiness

Get a demo →

holisticai.com/demo

Classify one system

The risk calculator walks the same sequence the Act does and names the provision behind each answer.

euaiact.com/eu-ai-act-risk-calculator

Check your exposure

The readiness check scores the whole estate against the obligations that already bind and the ones that are coming.

euaiact.com/eu-ai-act-readiness



The EU AI Act, as a working tool

euaiact.com is maintained by Holistic AI as an independent, source-verified reference: the complete legal text with the Omnibus amendments marked, a scope finder, a risk classifier, eight sector guides and a 31-step compliance checklist — all free.

YOUR NEXT STEP

Find out where you actually stand

Classify one system in a few minutes, or have Holistic AI assess the whole portfolio and produce the evidence behind it.

[Classify a system →](#)

[Start the assessment →](#)

euaiact.com/eu-ai-act-risk-calculator

holisticai.com/eu-ai-act-readiness

Or book a walkthrough of the platform: holisticai.com/demo