Standard Intelligence

White paper

Drawing the Line

Article 5 prohibited practices and where the boundaries actually sit

PUBLISHED

February 2025

WEB

standardintelligence.com

AUDIENCE

For AI governance leads, legal counsel and product compliance teams

CONTENTS

Summary 3

1. Article 5 bans practices outright, and no conformity process cures a caught one 4

2. Exposure turns on what a system can foreseeably do, and on what was done to prevent it 5

3. Each prohibition draws its boundary in a different place 6

- 3.1 Point (a): manipulation is prohibited by its effect, so engagement optimisation needs testing 7
- 3.2 Point (b): exploiting a vulnerability turns on whether the system protects or uses it 8
- 3.3 Point (c): social scoring reaches private actors, and the line falls at unrelated data and disproportionate treatment 8
- 3.4 Point (d): crime-risk prediction can reach beyond the police, including anti-money laundering screening 9
- 3.5 Point (e): facial scraping is banned when it builds or expands a recognition database 10
- 3.6 Point (f): emotion inference at work and in education is banned where it rests on biometric data 10
- 3.7 Point (g): biometric categorisation is banned for six sensitive attributes, including as a by-product 11
- 3.8 Point (h): real-time remote biometric identification is banned only for law enforcement, and its exceptions carry conditions 11

4. From 2 December 2026, generating non-consensual intimate imagery and child sexual abuse material is prohibited 12

5. One system can engage several prohibitions, and other law applies alongside them 13

6. Enforcement is national by default, and challenge can come from anyone affected 14

7. Screen the portfolio, record each conclusion with its conditions, and screen again when anything changes 15

8. The Standard Intelligence platform is designed to record the screen before a system goes live 16

Where to start, and what it costs 16

Appendix A: Article 5 quick reference 17

Appendix B: Regulatory timeline 18

Appendix C: Key terms and abbreviations 19

Basis and Standing 19

Sources 20

Summary

Article 5 is the one part of the AI Act where no documentation or conformity process can make a caught practice lawful. Its eight prohibitions in points (a) to (h) have applied since 2 February 2025, and two more, added by Regulation (EU) 2026/1744 in points (ba) and (bb), apply from 2 December 2026. The only routes out are the exceptions written into the text, and infringement sits in the Act's highest penalty tier: up to EUR 35 million or 7% of total worldwide annual turnover, whichever is higher.

For the manipulation and vulnerability bans, exposure does not depend on intent; it turns largely on what a system can foreseeably do, and on what was done to prevent it. Points (a) and (b) apply where a technique has "the objective, or the effect" of materially distorting behaviour, and the Commission's Guidelines say intent is sufficient but not necessary. The questions they treat as particularly relevant are whether the provider or deployer could reasonably have foreseen the significant harm and whether it took appropriate preventive and mitigating measures. The Guidelines apply similar reasoning to providers across all the prohibitions, and intent still bears on the size of any fine.

Each prohibition draws its boundary in a different place, and three are hardest to read. Point (a) separates prohibited manipulation from lawful persuasion, which puts engagement-optimising systems close to the line. Point (c) on social scoring reaches private as well as public actors, with the line falling at data from unrelated contexts and treatment that is unjustified or disproportionate. Point (h) bans real-time remote biometric identification in publicly accessible spaces only for law enforcement, under exceptions that carry their own conditions.

Two boundaries sit in a different place from where many readers assume. Point (d) on crime-risk prediction is not confined to the police: on the Commission's reading it can reach a bank's anti-money laundering screening, which falls outside it where the system uses the data the legislation specifies and its predictions are subject to human assessment and verification. Point (f) on emotion inference at work and in education reaches recruitment, but the Commission limits it to inferences from biometric data.

From 2 December 2026, AI systems that generate intimate imagery of identifiable people without their consent, or child sexual abuse material, are prohibited. A provider is caught where that generation is the system's intended purpose, or where it is a reasonably foreseeable and reproducible outcome and the system lacks reasonable and adequate safeguards to prevent it. For providers of general-purpose image, video and audio generators, the text itself now makes safeguards decisive.

One system can engage several prohibitions, and the boundaries move as capabilities change. Each prohibition must be assessed separately, a system that was clearly outside a definition can drift into it, and clearing Article 5 does not settle the system's position under data protection, consumer or product liability law.

Challenge can come from an authority or from anyone affected. National market surveillance authorities enforce Article 5, and since 2 August 2026 the AI Office has done so, with fining powers, for a defined set of systems. Anyone may complain to an authority,

and the Commission considers that the prohibitions have direct effect in national courts. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source.

The practical response is a screen, a record and a repeat screen. Map every AI system against all ten prohibitions, send the systems near a boundary for formal legal assessment, and record each conclusion with the testing, safeguards and conditions it depends on. Screen again whenever a system's capability, deployment scope or the guidance changes.

The first step is an inventory of AI systems with a plain-language description of each. Generative systems need screening against points (ba) and (bb) before 2 December 2026.

This paper states the position as at 11 October 2026. Basis and Standing, at the end, sets out which propositions are binding law, which are Commission guidance, and which are our own recommended controls.

1. Article 5 bans practices outright, and no conformity process cures a caught one

The AI Act entered into force on 1 August 2024, and its Chapter II prohibitions in Article 5 have applied since 2 February 2025 (Article 113, third paragraph, point (a)). They apply to systems placed on the market before that date as well as after, and the transitional rule for existing high-risk systems in Article 111(2) is expressly "without prejudice to the application of Article 5". Any provider or deployer that places on the market, puts into service or uses a system within one of the prohibitions is in breach of directly applicable EU law.

The Act sorts AI systems by level of risk, and Article 5 sits at the top. The Commission describes four levels: unacceptable, high, transparency (or limited) and minimal or no risk. High-risk systems carry extensive requirements because the legislature judged that safeguards can make them acceptable; the practices in Article 5 are treated as incapable of being made acceptable, so the prohibition is categorical rather than conditional.

There are now ten prohibitions, applying from two dates. Points (a) to (h) of Article 5(1) have applied since 2 February 2025. Regulation (EU) 2026/1744 (the Digital Omnibus on AI), in force since 27 July 2026, added points (ba) and (bb) and new paragraphs 1a and 1b, which apply from 2 December 2026 (Article 113, third paragraph, point (a)). Section 4 covers them.

The only routes out are the ones written into the text. Point (d) excludes systems that support a human assessment already based on objective and verifiable facts; point (f) excepts use for medical or safety reasons; point (g) excludes labelling or filtering of lawfully acquired biometric datasets and categorisation in law enforcement; and point

(h) permits three objectives on the conditions in paragraphs 2 to 7. Points (ba) and (bb) carry their own limits (section 4), and no other compliance pathway exists.

The prohibitions attach at three points in a system's life. Each applies to placing on the market, putting into service and use, except point (h), which covers use only; points (d) to (g) apply to putting into service "for this specific purpose". Providers and deployers are each caught within their own responsibilities, and section 2 explains why a provider with no contact with end users can still be exposed.

Infringement carries the highest fines in the Act. Article 99(3) sets fines of up to EUR 35 million or 7% of total worldwide annual turnover for the preceding financial year, whichever is higher; for SMEs, including start-ups, the cap is whichever is lower (Article 99(6)). Penalties have applied since 2 August 2025 (Article 113, third paragraph, point (b)). Article 99 sets maxima, and in setting a fine an authority is to have regard, as appropriate, to fines other market surveillance authorities have already imposed on the same operator for the same infringement (Article 99(7)(b)).

Fines are not the only consequence. Market surveillance authorities can require corrective action and can withdraw or recall a non-compliant system (Article 79 and Regulation (EU) 2019/1020).

Article 5 was made to apply first. The high-risk obligations now apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems (Article 113, as amended by Regulation (EU) 2026/1744), and a high-risk system already on the market before its date is caught only if its design changes significantly (Article 111(2), as amended). Article 5 has no such deferral or grandfathering, and providers had roughly six months from entry into force to screen their portfolios, which in our view was very little time for systems close to a boundary.

Several key terms are defined elsewhere in the Act. "Biometric data" (Article 3(34)), "emotion recognition system" (Article 3(39)), "real-time remote biometric identification system" (Article 3(42)) and "publicly accessible space" (Article 3(44)) all shape the prohibitions. The AI Act's definition of biometric data, unlike the GDPR's, does not require that the data allow unique identification, so analytics that never identify anyone can still rest on biometric data.

2. Exposure turns on what a system can foreseeably do, and on what was done to prevent it

Intent is sufficient for the manipulation and vulnerability bans, but not necessary. Points (a) and (b) apply to a technique with "the objective, or the effect" of materially distorting behaviour. The Commission's Guidelines say the intent of the provider or deployer is "a sufficient, but not a necessary, condition" (paragraph 83), and that point (a) covers "AI systems that manipulate individuals without any human intending them to

do so" (paragraph 69), for example where a system learns manipulative techniques from its training data.

What the Guidelines treat as particularly relevant is foreseeability and mitigation. In assessing whether significant harm is reasonably likely, they ask "whether the provider or deployer of the AI system could have reasonably foreseen the significant harm" and "whether they implemented appropriate preventive and mitigating measures" (paragraph 94). Manipulative behaviour that is merely incidental is not "purposeful" as long as the provider took such measures (paragraph 69), and harm caused by factors outside the system that the operator could not foresee is left out of the assessment (paragraph 96).

Similar logic shapes what the Guidelines expect of providers across all the prohibitions. Providers should not place on the market systems, including general-purpose AI systems, that are "reasonably likely to behave or be directly used in a manner prohibited by Article 5 AI Act", and are expected to build in safeguards against reasonably foreseeable misuse (paragraph 40). Deployers must not use a system in a prohibited way, "including not to bypass any safety guardrails implemented by the providers of the system" (paragraph 40), and a provider's terms of use excluding a prohibited use do not protect the deployer (paragraph 14).

A deployer of a bought-in system cannot rely on its own lack of intent. If prohibited effects occur in its deployment and readily available information about the system's capabilities made them foreseeable, the deployer is exposed even though it did not design the system. Intent still matters to the consequences: an authority setting a fine is to have regard, as appropriate, to "the intentional or negligent character of the infringement" and to the technical and organisational measures the operator implemented (Article 99(7)(i) and (g)).

"Significant harm" is not defined in the Act. Points (a) and (b) turn on it, and the Guidelines offer factors (severity, context and cumulative effects, scale, the vulnerability of those affected, and duration and reversibility) without a threshold (paragraphs 85 to 94). The Guidelines are non-binding, so the threshold will be settled by enforcement decisions and, in the end, by the Court of Justice.

The best evidence of foreseeability and mitigation is a record of testing and safeguards. An organisation that can show what it foresaw, what it tested and what it put in place to prevent a prohibited effect can answer the questions the Guidelines put in paragraph 94. Section 7 sets out that record.

3. Each prohibition draws its boundary in a different place

Points (a), (c) and (h) raise the hardest interpretive questions, and points (d) and (f) draw theirs in a different place from where many readers assume.

3.1 Point (a): manipulation is prohibited by its effect, so engagement optimisation needs testing

The prohibition. Point (a) bans AI systems that deploy subliminal techniques beyond a person's consciousness, or purposefully manipulative or deceptive techniques, with the objective or effect of materially distorting a person's or group's behaviour by appreciably impairing their ability to make an informed decision, causing them to take a decision they would not otherwise have taken, in a manner that causes or is reasonably likely to cause significant harm.

It has two limbs, and both need distortion and harm. The first limb covers techniques operating "beyond a person's consciousness", the classical subliminal case; the second, wider limb covers "purposefully manipulative or deceptive techniques". Both require material distortion of behaviour, which the Guidelines describe as going beyond lawful persuasion (paragraph 78), and actual or reasonably likely significant harm.

Awareness of the influence is not always an answer. The Guidelines, drawing on Recital 29, say the prohibition also covers techniques that people may be unable to control or resist even when they are aware of the attempt to influence them (paragraph 68).

"Purposefully" does not require a purpose held by a person. The Guidelines describe purposefully manipulative techniques as those "designed or objectively aim to influence, alter, or control an individual's behaviour in a manner that undermines their individual autonomy and free choices", and say intent is not necessary (paragraphs 67, 69 and 76). A system trained to maximise engagement by exploiting variable-ratio reinforcement could therefore be treated as purposefully manipulative even if the developer's conscious goal was only to retain users.

The line runs between manipulation and persuasion. The Guidelines do not address engagement-maximising systems by name, but they say that a system using "personalised recommendations based on transparent algorithms and user preferences and controls engages in persuasion" (paragraph 128). An engagement objective pursued through manipulative techniques is exposed if the other conditions are met; one pursued through transparent recommendations with user control sits on the persuasion side.

A generic disclosure does not move a system across the line. The Guidelines list transparency among the measures that help keep a system outside the prohibition (paragraph 95), but a privacy notice saying the service uses "advanced machine learning to personalise your experience" does not tell the user about any specific mechanism. If the techniques are manipulative in the point (a) sense, a notice of that kind does not displace the prohibition.

What follows. Providers and deployers of recommendation engines, personalisation systems, engagement optimisation tools and consumer-facing AI interfaces should review the influence mechanisms their systems use, starting with systems whose training objectives are defined by engagement metrics with no constraint on mechanism. Because the prohibition can be engaged by effect, providers should test, for example through A/B testing, whether behaviour-distorting effects emerge from training, and keep the results (section 7).

3.2 Point (b): exploiting a vulnerability turns on whether the system protects or uses it

The prohibition. Point (b) bans AI systems that exploit any vulnerability of a natural person or a specific group due to their age, disability or a specific social or economic situation, with the objective or effect of materially distorting the behaviour of that person or a member of that group in a manner that causes or is reasonably likely to cause significant harm.

The three sources of vulnerability are a closed list. The Guidelines say that exploiting the vulnerabilities of groups other than those defined by age, disability or a specific social or economic situation falls outside point (b) (paragraphs 102 and 103). The third category is broader than it looks: it can include transient circumstances such as temporary unemployment or over-indebtedness, though not loneliness or grievances that anyone may experience (paragraph 109).

The test is whether the system is calibrated to protect the vulnerability or to use it. A system that detects signs of financial distress and routes the user to debt advice is not exploiting that vulnerability; one that detects the same signs and gives high-margin credit products more prominence may be. The prohibition can also be engaged indirectly, through people who are not the system's users. An employer's system that identifies staff likely to be in financial distress and adjusts their working conditions could fall within it if it materially distorts their behaviour in a way that causes or is reasonably likely to cause significant harm.

What follows. Deployers of consumer-facing systems should establish whether their users include people in the three categories and, if so, whether the system's behaviour toward them protects or exploits the vulnerability. The exposure is greatest in health, financial services, social care and digital platforms.

3.3 Point (c): social scoring reaches private actors, and the line falls at unrelated data and disproportionate treatment

The prohibition. Point (c) bans AI systems that evaluate or classify natural persons or groups over a certain period of time based on their social behaviour or known, inferred or predicted personal or personality characteristics, where the resulting social score leads to detrimental or unfavourable treatment (i) in social contexts unrelated to those in which the data was originally generated or collected, or (ii) that is unjustified or disproportionate to their social behaviour or its gravity.

Three elements must be present together. They are evaluation or classification over a period of time; a basis in social behaviour or personal or personality characteristics; and detrimental or unfavourable treatment either in an unrelated context or that is unjustified or disproportionate. The two forms of treatment are alternatives, and "unfavourable" treatment need not involve harm (Guidelines, paragraphs 163 to 165).

The ban is not confined to the state. The Commission's 2021 proposal limited it to public authorities; the final text removed that limitation, and the Guidelines read it as covering "both the public and the private sector" (paragraph 152). Large-scale private scoring, by credit reference agencies, gig economy platforms, insurance telematics providers and social media platforms, can fall within point (c) where it produces a social score leading to detrimental treatment in unrelated contexts or out of proportion to the behaviour assessed.

Lawful evaluation for a specific purpose sits outside it, on conditions. The Guidelines give credit scoring based on relevant financial data, fraud evaluation based on relevant data, and telematics-based insurance pricing where the premium increase is proportionate to the risky driving, as examples outside the prohibition (paragraph 177). Each example carries a condition about relevant data or proportionate treatment, and practices outside point (c) remain subject to the GDPR and sector law.

What follows. Organisations running continuous behavioural scoring should map its data inputs, its method and, above all, every downstream use of the score. Scoring whose effects stay within the context in which the data was collected is at lower risk. Any organisation, public or private, running a scoring system whose output crosses contexts should treat point (c) as a live risk until the analysis shows otherwise.

3.4 Point (d): crime-risk prediction can reach beyond the police, including anti-money laundering screening

The prohibition. Point (d) bans AI systems that make risk assessments of natural persons to assess or predict the risk of a person committing a criminal offence, based solely on profiling or on assessing their personality traits and characteristics. It does not apply to systems that support the human assessment of a person's involvement in a criminal activity "which is already based on objective and verifiable facts directly linked to a criminal activity", and administrative offences are outside its scope (Guidelines, paragraph 217).

It targets individual prediction from profiles. The prohibition does not reach all AI-assisted criminal risk tools, only those producing individual risk assessments from group profiles or personality inference without grounding in facts about the individual. Recidivism and reoffending risk tools are directly implicated to the extent that they rely solely on profiling or personality traits.

The exclusion for existing objective evidence is narrow. The facts must be objective, verifiable, directly linked to criminal activity and already established, and the system must support rather than replace a human assessment. If an officer is shown an AI risk score in a setting where overrides rarely occur, the review may be a rubber stamp, and the system is then arguably not supporting a human assessment in any meaningful sense. Departure rate, the proportion of decisions in which the person with authority disagreed with the system's output and acted to change it, gives a first indication of whether the review leaves any trace, though not proof either way.

The prohibition is not confined to law enforcement. The Commission's Guidelines say it may apply to a private entity where assessing crime risk is objectively necessary to comply with a legal obligation, such as a bank's anti-money laundering screening, and that such use falls outside it where the system uses only the data that legislation specifies and its predictions are subject to human assessment and verification in compliance with that legislation (paragraph 209). Profiling by a private entity for its own business, safety or financial interests, without the purpose of predicting a specific criminal offence, is not covered (paragraphs 210 and 211).

What follows. Law enforcement and criminal justice agencies, and private entities under a legal duty to assess crime risk such as banks screening for money laundering, should assess whether each tool meets the exclusion in point (d) or the paragraph 209 conditions, and whether human assessment is substantive rather than formal. Treating a

system as outside the prohibition rather than within it is among the most consequential classifications an organisation will make.

3.5 Point (e): facial scraping is banned when it builds or expands a recognition database

The prohibition. Point (e) bans AI systems that create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage.

"Untargeted" means without a focus on a given individual or group. The Guidelines describe untargeted scraping as operating like a "vacuum cleaner" (paragraph 228); scraping for images of a specific suspect or a defined group of victims is targeted and outside the ban. Building a database step by step, selecting a different group each time, is caught where the end result is functionally the same as untargeted scraping (paragraph 230).

Expanding an existing database is caught as well as creating one. A company that uses AI to add scraped images to an established database is engaged in prohibited activity even if the database predates 2 February 2025; a database built before that date and not further expanded in this way must still comply with data protection law (paragraph 236). The explicit reference to CCTV footage reaches systems that extract and aggregate facial images from such footage.

Some uses fall outside. The ban does not cover scraping of other biometric data such as voice samples, or facial image databases used to train or test models where the persons are not identified (paragraph 234).

What follows. Providers of facial recognition systems should audit the provenance of their recognition databases. Deployers using third-party facial recognition services should require contractual representations about how the database underpinning the service was built and expanded.

3.6 Point (f): emotion inference at work and in education is banned where it rests on biometric data

The prohibition. Point (f) bans AI systems used to infer the emotions of a natural person in the workplace and in education institutions, except where the use is intended to be put in place or into the market for medical or safety reasons.

The Commission confines the ban to inference from biometric data. Point (f) does not use the defined term "emotion recognition system", but the Guidelines read it with that definition (Article 3(39), which covers emotions or intentions inferred "on the basis of their biometric data") and limit it "to inferences based on a person's biometric data" (paragraph 245). Sentiment inferred from written text is therefore outside the ban (paragraph 251), though it may raise data protection questions. Biometric emotion recognition outside the ban is high-risk under Annex III point 1(c).

"Workplace" is read broadly and includes recruitment. The Guidelines treat the workplace as any physical or virtual space where people carry out assigned work, give hybrid video calls as an example, and apply the ban to candidates during selection and hiring (paragraph 254). Video interview tools that infer candidates' emotional states from

facial expressions therefore engage it, and in our reading so does inference from an employee's camera while working from home.

The exception is narrow and turns on purpose. The Guidelines say it "should be narrowly interpreted" (paragraph 256): general monitoring of stress or wellbeing is not covered (paragraph 257), and safety means protection of life and health, not of property against theft or fraud (paragraph 258). Detecting fatigue is not emotion recognition at all (Recital 18).

What follows. Providers and deployers of systems used in employment and education should check whether any feature infers emotions or intentions from biometric data. Features that look peripheral (attention tracking, engagement scoring, sentiment analysis from video) can engage the ban, and reliance on the exception needs a documented medical or safety purpose of the kind the Guidelines accept.

3.7 Point (g): biometric categorisation is banned for six sensitive attributes, including as a by-product

The prohibition. Point (g) bans biometric categorisation systems that categorise natural persons individually on their biometric data to deduce or infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation. It does not cover labelling or filtering of lawfully acquired biometric datasets, such as images, or categorising of biometric data in the area of law enforcement.

The ban is aimed at the inference itself. It targets systems that use biometric inputs (facial features, gait, voice, physiological measurements) to produce outputs corresponding to the six attributes, whether or not the inference is accurate.

Intermediate inferences are an open question. It is not settled whether a system that uses a sensitive-attribute inference as an intermediate step in a permitted process is within the ban. A facial recognition system that uses race-correlated features in identifying a specific known person, without categorising anyone by race, may or may not engage it depending on how the computation is characterised.

What follows. Any system processing biometric data should be checked for whether its training data, architecture or outputs include categorisation corresponding to the six attributes, because systems built for legitimate purposes can produce prohibited inferences as by-products. Vendors should document in model cards which inferences their systems can and cannot produce.

3.8 Point (h): real-time remote biometric identification is banned only for law enforcement, and its exceptions carry conditions

The prohibition. Point (h) bans the use of real-time remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement, unless and in so far as such use is strictly necessary for one of three objectives, and then only on the conditions in Article 5(2) to (7), including national legislation permitting such use (Article 5(5)).

It covers use, not development or sale. Point (h) prohibits only a specific use: in publicly accessible spaces, for law enforcement, in real time. It has the most elaborate exception framework in Article 5, permitting use in defined circumstances under prior authorisation.

Non-law enforcement use is outside it. The same system that a police force may not use is not prohibited under Article 5 when a private operator uses it in a publicly accessible space for another purpose. Such processing of biometric data remains subject to Article 9 of the GDPR, and remote biometric identification systems are high-risk under Annex III point 1(a).

The three objectives are narrow. They are the targeted search for specific victims of abduction, trafficking or sexual exploitation, and for missing persons; the prevention of a specific, substantial and imminent threat to life or physical safety, or a genuine and present or genuine and foreseeable threat of a terrorist attack; and the localisation or identification of a person suspected of an offence listed in Annex II that is punishable in the Member State concerned by a custodial sentence or detention order of at least four years.

Each use needs authorisation, an impact assessment and notification. The authorisation must come in advance from a judicial authority or an independent administrative authority whose decision is binding; in a duly justified situation of urgency, use may begin without it if authorisation is requested without undue delay, "at the latest within 24 hours" (Article 5(3)). The law enforcement authority must have completed a fundamental rights impact assessment and registered the system in the EU database, or in urgent cases complete registration without undue delay (Article 5(2)), and must notify each use to the market surveillance authority and the data protection authority (Article 5(4)).

No adverse legal decision may rest on the match alone. Article 5(3) provides that "No decision that produces an adverse legal effect on a person may be taken based solely on the output" of the system.

What follows. Where a Member State has legislated under Article 5(5), law enforcement authorities need authorisation processes meeting Article 5(3) and notification processes meeting Article 5(4). Vendors supplying these systems to law enforcement should reflect their customers' obligations in procurement terms, specifications and use policies, and deployers should keep authorisation, exception, scope and notification records in a governed, version-controlled register.

4. From 2 December 2026, generating non-consensual intimate imagery and child sexual abuse material is prohibited

Two prohibitions apply from 2 December 2026. Point (ba) bans AI systems that generate or manipulate realistic images, videos, audio or similar material of an identifiable person's intimate parts, or of an identifiable person engaged in sexually explicit activities, without that person's "freely-given, specific, informed, unambiguous and explicit consent". Point (bb) bans AI systems that generate or manipulate child sexual abuse material or performance within the meaning of Article 2, points (c) and (e), of Directive 2011/93/EU, "except where a 'without right' defence applies under national law".

Paragraph 1a limits when providers and deployers are caught. Placing on the market or putting into service is prohibited only where the generation or manipulation is the system's intended purpose, or where its "design, training, architecture, capabilities or user-facing functionalities" make it a "reasonably foreseeable and reproducible outcome, without requiring significant technical modification", and the system lacks "reasonable and adequate technical safety measures and other safeguards to reliably prevent" it, taking into account reasonably foreseeable misuse, and to correct observed or reported misuse (Article 5(1a)(a)). Use is prohibited only where the deployer uses the system for the purpose of generating or manipulating such material (Article 5(1a)(b)).

Paragraph 1b narrows "manipulation" for point (ba). Altering material in a way that does not increase the exposure of any depicted intimate parts, or alter the nature of any depicted sexually explicit activity, is not manipulation for this purpose (Article 5(1b)).

The provider limb puts the foreseeability test into the text. For these two points, a provider of a general-purpose image, video or audio generator is caught not only where generation is the intended purpose but also by what its system can reproducibly do and whether its safeguards reliably prevent it, the same reasoning section 2 draws from the Guidelines for points (a) and (b). The safeguards must also correct misuse once it is observed or reported, so monitoring and response are part of the test.

What follows. Providers of generative systems that can produce realistic images, video or audio of people should, before 2 December 2026, test whether the prohibited outputs are reasonably foreseeable and reproducible, put in place safeguards that reliably prevent them, and set up a route for reporting and correcting misuse. Organisations that integrate such systems into their own services should check those safeguards, and deployers should ensure no use of the system is directed at generating the material.

5. One system can engage several prohibitions, and other law applies alongside them

Prohibitions combine, so each must be assessed separately. A consumer platform that personalises content using manipulation calibrated to identified vulnerabilities, and produces scores used across different service domains, could engage points (a), (b) and (c) at once. If it also uses facial recognition in its shops, point (h) does not apply, because it is limited to law enforcement. The system engages point (e) or (g) only if it involves untargeted scraping or sensitive-attribute categorisation, and point (f) if it infers the emotions of staff from biometric data.

The boundaries move as capabilities change. The prohibitions rest on concepts ("subliminal", "emotion", "biometric", "real-time") that are stable at their core but tested at their edges by new technology. A system clearly outside a definition when it was assessed may, as its capabilities develop or its deployment widens, produce effects that bring it within one, so an assessment should consider reasonably foreseeable developments over the deployment period, not only current capability.

General-purpose AI is reached through the systems built on it. The prohibitions apply to any AI system, whether it has an intended purpose or is general-purpose (Guidelines, paragraph 39), and the expectations on providers in paragraph 40 extend to general-purpose AI systems. How far the provider of a general-purpose AI model answers for prohibited systems that others build on it is not resolved by the Act's text, though points (ba) and (bb) now address part of that question for generative systems (section 4).

Clearing Article 5 does not settle other law. Article 5(8) states that the Article does not affect prohibitions that apply where an AI practice infringes other Union law, and the Guidelines say that respecting the prohibitions is "not a sufficient condition" for compliance with other legislation (paragraph 43). Exposure under Article 5 often occurs alongside exposure under the GDPR (Article 22 on solely automated decisions with legal or similarly significant effects, and Article 9 on special category data, including biometric data processed to identify a person uniquely), the Law Enforcement Directive and the Unfair Commercial Practices Directive.

Product liability is the other parallel regime. The revised Product Liability Directive, Directive (EU) 2024/2853, brings software, and so AI systems, within the definition of a product (Article 4(1)) for products placed on the market or put into service after 8 December 2026 (Article 2(1), as corrected). The separate AI Liability Directive proposal was withdrawn on 6 October 2025.

6. Enforcement is national by default, and challenge can come from anyone affected

National market surveillance authorities enforce Article 5 in their territories. Each Member State decides which of its authorities supervises the prohibitions, and some have divided them by practice. The European Data Protection Board has no enforcement role under the AI Act.

The AI Office enforces it for a defined set of systems. Since 2 August 2026 the AI Office has been exclusively competent for AI systems based on a general-purpose AI model where the model and the system come from the same provider or undertaking, subject to exceptions, and for AI systems that constitute or are integrated into a very large online platform or search engine (Article 75(1), as amended by Regulation (EU) 2026/1744). For those systems it can impose fines on the Article 99 scale, including the Article 5 tier (Article 75c(4)).

Early enforcement is most likely where authorities are already in place. In our view, the operationally advanced Member States, Spain, Italy, Ireland and Denmark, may move first, while France and the Netherlands had not completed their designations as at 11 October 2026 (see our paper *Enforcement Architecture*, section 3).

Challenge does not depend on an authority acting first. Any natural or legal person with grounds to consider that the Act has been infringed may complain to the relevant market surveillance authority (Article 85). The Commission considers that the

prohibitions “have direct effect and thus enable affected parties to enforce them in national courts and request interim injunctions” (Guidelines, paragraph 431).

Reporting persons are protected. Article 87 applies Directive (EU) 2019/1937 to reports of AI Act infringements and the protection of those who make them, from 2 August 2026, so an employee who reports a suspected Article 5 breach to a competent authority is protected on the conditions that Directive sets.

No fine had been confirmed as at the position date. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source.

7. Screen the portfolio, record each conclusion with its conditions, and screen again when anything changes

Every organisation that develops, deploys or uses AI in the EU should complete an Article 5 portfolio screen. It maps each AI system against all ten prohibitions, identifies those close to a boundary and sends them for deeper legal and technical assessment. Systems clearly within points (a) to (h) should have been withdrawn before 2 February 2025, and systems clearly within points (ba) or (bb) need to be withdrawn or brought outside them before 2 December 2026.

The screen has three steps:

(a) produce a plain-language description of each system’s intended purpose, deployment context and main user population;
(b) send each system close to a boundary for formal legal assessment rather than an informal view, starting with recommendation, personalisation, biometric processing, behavioural scoring and generative systems;
(c) record every system, whether assessed, classified or pending review, in a central AI governance register.

Each Article 5 conclusion should sit in a governed, version-controlled record. The record should hold the system description, the classification and compliance position, the testing done and the safeguards relied on (section 2), and any conditions on which lawful operation depends. This is our recommended control, not a requirement of the Act. It is, however, the evidence an authority or a court is likely to ask for, and a record assembled under the pressure of an investigation cannot reliably reconstruct what was foreseen and tested at the time.

Each conclusion should be revisited when its basis changes. That means a change in a system’s capability, training or deployment scope; a new use of a score or output; and new Commission guidance, a harmonised standard or an enforcement decision. A trigger that flags these changes as they happen is more reliable than periodic manual review.

8. The Standard Intelligence platform is designed to record the screen before a system goes live

The practical risk under Article 5 is deploying close to a boundary without having tested where it sits. The Standard Intelligence platform is designed to surface prohibition risk and record the assessment before a system goes live.

Standard Intelligence platform: Article 5 screening capability map

Platform capability	How it supports Article 5 screening
AI Regulatory Navigator	Screens a described system against the Article 5 prohibitions and returns a cited assessment of where it sits relative to each boundary.
AI Intelligence Centre	Records the screening decision and its rationale, dated before deployment, as evidence that the boundary was assessed.
Approval Workflow	Holds go-live behind a named approval, so a system that touches a prohibition boundary cannot be deployed without explicit sign-off.

A governance lead can filter the register by prohibition to find systems engaging more than one point, and continuous gap analysis is designed to flag when a change in a system's deployment scope creates new Article 5 exposure.

Where to start, and what it costs

The work comes in stages, and the first uses information the organisation already holds.

This month: the inventory and first screen. Describe each AI system in plain language and map it against the ten prohibitions (section 7). This finds the systems that need more than a first look.

Before 2 December 2026: the generative systems. Test any system that can generate realistic images, video or audio of people against points (ba) and (bb), and confirm that its safeguards reliably prevent the prohibited outputs and that misuse can be reported and corrected (section 4).

A quarter, with legal: the systems near a boundary. Give each one a formal assessment and a record of the testing, safeguards and conditions on which its position depends (sections 2 and 7). The heaviest cost sits here, in legal and technical time on the minority of systems close to a line.

Ongoing: screen again on change. Treat a change in capability, deployment scope, use of outputs or guidance as a trigger for a fresh screen.

On timing. Points (a) to (h) already apply, penalties have applied since 2 August 2025, and the Commission considers the prohibitions enforceable by affected persons in national courts. As at 11 October 2026 no fine under the AI Act had been confirmed by any

regulator or official source. No conformity process can cure a caught practice, however, and existing systems are not grandfathered, so waiting does not reduce the exposure.

Deferring has three costs:

- a system found within a prohibition has to stop, whatever has been invested in it;
- the record of what was foreseen and tested is assembled after the event, when it is least persuasive;
- generative systems that reach 2 December 2026 without tested safeguards fall within points (ba) and (bb) from that date.

Article 5 cannot be satisfied by documentation, but an organisation that has screened its portfolio and kept the record can say, system by system, where each one sits relative to each line and what keeps it there.

Appendix A: Article 5 quick reference

Provision	Prohibited practice	Key conditions	Exceptions and limits in the text
5(1)(a)	Subliminal, manipulative or deceptive techniques	Material distortion of behaviour; appreciably impaired informed decision; significant harm, actual or reasonably likely	None
5(1)(b)	Exploitation of vulnerabilities	Age, disability, specific social or economic situation; material distortion; significant harm	None
5(1)(ba) (from 2 December 2026)	Non-consensual intimate imagery of identifiable persons	Realistic generated or manipulated material; no freely given, specific, informed, unambiguous and explicit consent	Paragraph 1a limits on providers and deployers; paragraph 1b on manipulation
5(1)(bb) (from 2 December 2026)	Child sexual abuse material	Material or performance within Directive 2011/93/EU, Article 2(c) and (e)	“Without right” defence under national law; paragraph 1a limits
5(1)(c)	Social scoring	Evaluation over time; detrimental or unfavourable treatment, cross-contextual or unjustified or disproportionate	None
5(1)(d)	Crime-risk prediction	Criminal offences only; based solely on profiling or personality traits	Support for a human assessment already based on objective and verifiable facts directly linked to criminal activity

Provision	Prohibited practice	Key conditions	Exceptions and limits in the text
5(1)(e)	Facial image database scraping	Untargeted; from the internet or CCTV footage; creating or expanding a recognition database	None
5(1)(f)	Emotion inference at work and in education	Workplace (including recruitment); education institutions; biometric data (Commission reading)	Medical or safety reasons
5(1)(g)	Biometric categorisation for sensitive attributes	Race, political opinions, trade union membership, religious or philosophical beliefs, sex life, sexual orientation	Labelling or filtering of lawfully acquired biometric datasets; categorising biometric data in law enforcement
5(1)(h)	Real-time remote biometric identification	Publicly accessible spaces; law enforcement purposes	Three objectives, with authorisation and conditions under paragraphs 2 to 7 and national law

Appendix B: Regulatory timeline

Date	Event
1 August 2024	AI Act enters into force
2 February 2025	Article 5(1)(a) to (h) apply
4 February 2025	Commission approves the content of its Guidelines on prohibited AI practices (C(2025) 884 final)
29 July 2025	Guidelines adopted as C(2025) 5052 final
2 August 2025	Penalties apply; general-purpose AI model obligations apply; Member States to designate authorities and publish contact details (Article 70(2))
27 July 2026	Regulation (EU) 2026/1744 enters into force
2 August 2026	General date of application, including the AI Office's exclusive competence under Article 75(1) and Article 87
2 December 2026	Article 5(1)(ba) and (bb), and paragraphs 1a and 1b, apply
8 December 2026	Revised Product Liability Directive applies to products placed on the market or put into service after this date
2 August 2027	General-purpose AI models placed on the market before 2 August 2025 must comply (Article 111(3))
2 December 2027	Annex III high-risk obligations apply
2 August 2028	Annex I high-risk obligations apply

Appendix C: Key terms and abbreviations

Statutory terms cite Article 3 of the AI Act.

Biometric data. Personal data resulting from specific technical processing relating to the physical, physiological or behavioural characteristics of a natural person, such as facial images or dactyloscopic data (Article 3(34)).

Emotion recognition system. An AI system for the purpose of identifying or inferring emotions or intentions of natural persons on the basis of their biometric data (Article 3(39)).

Publicly accessible space. Any publicly or privately owned physical place accessible to an undetermined number of natural persons, regardless of whether certain conditions for access may apply, and regardless of the potential capacity restrictions (Article 3(44)).

Purposefully manipulative techniques. The second limb of point (a); the Commission reads it as not requiring intent (Guidelines, paragraphs 67, 69 and 76).

Real-time remote biometric identification system. A remote biometric identification system (an AI system for identifying natural persons, without their active involvement, typically at a distance, by comparing their biometric data with a reference database: Article 3(41)) in which capture, comparison and identification all occur without a significant delay (Article 3(42)).

Departure rate. The proportion of decisions in which the person with authority disagreed with the system's output and acted to change it (a Standard Intelligence term).

Abbreviations.

- GDPR: General Data Protection Regulation, Regulation (EU) 2016/679
- SME: small and medium-sized enterprise

This paper describes the position as at 11 October 2026 and is not legal advice. The Commission's Guidelines on prohibited AI practices are non-binding, and authoritative interpretation of Article 5 is for the Court of Justice. Readers should take qualified advice before acting.

Basis and Standing

Binding law.

- The AI Act, Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744 (in force 27 July 2026): Articles 3(34), (39), (41), (42) and (44); Article 5, including points (ba) and (bb) and paragraphs 1a and 1b inserted by Regulation (EU) 2026/1744; Articles 75(1) and 75c(4) as amended or inserted; Articles 70(2), 79, 85, 87, 99(3), (6) and (7) and

111(2) and (3); Annex II; Annex III points 1(a) and 1(c); and Recitals 18 and 29, as aids to interpretation. Article 5(1)(a) to (h) have applied since 2 February 2025 and points (ba) and (bb) apply from 2 December 2026 (Article 113, third paragraph, point (a)); penalties since 2 August 2025 (point (b)); Annex III high-risk obligations from 2 December 2027 and Annex I from 2 August 2028 (point (c)), subject to Article 111(2).
- Directive 2011/93/EU, Article 2(c) and (e).
- The GDPR, Articles 9 and 22; the Law Enforcement Directive (EU) 2016/680; the Unfair Commercial Practices Directive 2005/29/EC; Directive (EU) 2019/1937; Regulation (EU) 2019/1020.
- Directive (EU) 2024/2853, Articles 2(1), as corrected by the corrigendum of 7 May 2026, and 4(1).

Commission guidance, not law. The Guidelines on prohibited AI practices, C(2025) 5052 final of 29 July 2025, are non-binding (paragraph 5), and their examples are indicative only. This paper relies on paragraphs 14, 39, 40, 43, 67 to 69, 76, 78, 83, 85 to 96, 102, 103, 109, 128, 152, 163 to 165, 177, 209 to 211, 217, 228, 230, 234, 236, 245, 251, 254, 256 to 258 and 431. The reading of point (f) as limited to biometric data, the paragraph 209 conditions for anti-money laundering screening and the treatment of foreseeability and mitigation are the Commission's views and untested in court. The four levels of risk are the Commission's description, not terms of the Act.

Not a regulatory requirement. The portfolio screen, the governed record of each conclusion, the change triggers, the testing recommended in sections 3.1 and 4, and the vendor representations and model cards in sections 3.5 and 3.7 are our recommended controls. They reflect the Guidelines' emphasis on foreseeability and preventive measures (paragraphs 40, 69 and 94) and, for points (ba) and (bb), the safeguards test in Article 5(1a)(a)(ii). Departure rate is a signal of whether a human assessment leaves any trace, not proof that it is substantive.

Our own views. The examples of engagement optimisation, financial distress targeting and employee distress detection, the reading of home working as within the workplace, the open question on intermediate inferences under point (g), and the forecast of which Member States may enforce first are our analysis.

Research. This paper relies on no empirical research.

Standing rule. A documented conclusion records a position; it does not make a prohibited practice lawful.

Enforcement position. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source.

Sources

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

2. Regulation (EU) 2024/1689 (AI Act), consolidated text of 27 July 2026: Articles 3, 5, 70, 75, 75c, 79, 85, 87, 99, 111 and 113; Annexes II and III; Recitals 18 and 29. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng
3. Regulation (EU) 2026/1744 of 8 July 2026 (Digital Omnibus on AI), OJ L, 24.7.2026, inserting Article 5(1)(ba) and (bb), (1a) and (1b) and Article 75c, and amending Articles 75, 111 and 113. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
4. European Commission, Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act), C(2025) 5052 final, 29 July 2025 (content approved 4 February 2025; non-binding). https://ai-act-service-desk.ec.europa.eu/sites/default/files/2026-01/guide-prohibited_en.pdf
5. European Commission, AI Act policy page (the four levels of risk). https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
6. European Commission, Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), COM(2021) 206 final, 21 April 2021. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206
7. Directive 2011/93/EU of the European Parliament and of the Council of 13 December 2011 on combating the sexual abuse and sexual exploitation of children and child pornography, Article 2. https://eur-lex.europa.eu/eli/dir/2011/93/oj
8. Regulation (EU) 2016/679 (General Data Protection Regulation), Articles 9 and 22. https://eur-lex.europa.eu/eli/reg/2016/679/oj
9. Directive (EU) 2016/680 (Law Enforcement Directive). https://eur-lex.europa.eu/eli/dir/2016/680/oj
10. Directive 2005/29/EC (Unfair Commercial Practices Directive). https://eur-lex.europa.eu/eli/dir/2005/29/oj
11. Directive (EU) 2019/1937 on the protection of persons who report breaches of Union law. https://eur-lex.europa.eu/eli/dir/2019/1937/oj
12. Regulation (EU) 2019/1020 on market surveillance and compliance of products. https://eur-lex.europa.eu/eli/reg/2019/1020/oj
13. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC, Articles 2(1) and 4(1). https://eur-lex.europa.eu/eli/dir/2024/2853/oj
14. Corrigendum to Directive (EU) 2024/2853, 7 May 2026 (Article 2(1) date). https://eur-lex.europa.eu/eli/dir/2024/2853/corrigendum/2026-05-07/oj
15. Withdrawal of Commission proposals, including the AI Liability Directive proposal (COM(2022) 496), OJ C/2025/5423, 6 October 2025. https://eur-lex.europa.eu/eli/C/2025/5423/oj
16. Standard Intelligence, *Enforcement Architecture: How national competent authorities will supervise AI Act compliance*, section 3. https://standardintelligence.com/whitepapers/enforcement-architecture

Provisions are cited by article and paragraph in the text above.

Standard Intelligence

GET IN TOUCH

Screen your deployments against the bans.

EMAIL

hello@standardintelligence.com

WEB

standardintelligence.com

LOCATION

Manchester, United Kingdom

AI Use: AI tools were used in research, editorial review and drafting support. All sources were retrieved and checked against the published record by the author, who holds editorial responsibility for the contents and for the conclusions drawn. This disclosure is made voluntarily. To the extent the Article 50(4) obligation in Regulation (EU) 2024/1689 applies to this paper as text published to inform the public on matters of public interest, the exception for content that has undergone human review and for which a person holds editorial responsibility applies.

STANDARD INTELLIGENCE LIMITED · REGISTERED IN ENGLAND AND WALES · COMPANY NO. 16737890