Prohibited AI Practices
AI uses completely banned

Co-authored with Markus Krebsz, **The Human AI Institute**, *Founder*

30 June 2025

2. Risk Classification of AI Systems — Levels, Impact, Categories, Compliance

2.1 Understanding the Risk-Based Approach
AI systems grouped by risk

2.2 Prohibited AI Practices
AI uses completely banned

2.3 High-Risk AI Systems: Criteria and Examples
Identifying regulated high-risk systems.

2.4 Obligations for High-Risk AI Systems
Rules for high-risk system compliance

Setting Red Lines: Prohibited AI Practices Under the AI Act

While the EU AI Act embraces a flexible, risk-based regulatory approach, it also draws a firm line when it comes to certain applications of artificial intelligence. Article 5 outlines eight prohibited practices, posing significant compliance risks for companies that are unaware of or inadvertently engage in them. These practices are deemed fundamentally incompatible with European values—such as respect for human dignity, autonomy, and fundamental rights. In such cases, the regulation does not seek to manage risk; it prohibits the use altogether.

These banned AI practices are classified under the **"unacceptable risk"** category, and they are considered so harmful that their use is illegal across the European Union, with very few exceptions. This zero-tolerance stance is designed to uphold the foundational principles of the EU legal order in the face of rapidly evolving technological capabilities.

Why Some AI Systems Are Banned Entirely

The EU's decision to ban certain AI uses reflects a key ethical and legal judgement: that not all innovation is acceptable, and that some uses of AI—regardless of their effectiveness or market demand—pose risks that cannot be justified. These prohibitions are not based solely on technical considerations, but on deep-seated societal concerns about surveillance, manipulation, discrimination, and exploitation.

The banned practices are narrowly defined to avoid sweeping restrictions, but their scope is significant. They include systems that threaten democratic participation, social equality, mental autonomy, or physical safety in ways the EU deems unacceptable in a rights-respecting society.

Overview of Prohibited AI Practices

The AI Act specifies eight main categories of prohibited AI systems. These systems are banned outright, with only tightly limited exceptions in some cases, particularly where public safety is at stake.

1. Manipulative or Deceptive AI (Art. 5.1.a)

Rationale and Objectives

Article 5(1)(a) of the AI Act prohibits AI systems that deploy subliminal, manipulative, or deceptive techniques that materially distort behavior and are likely to cause significant harm. The rationale behind this prohibition is the protection of individual autonomy, mental integrity, and dignity, as well as the preservation of informed decision-making in democratic societies. The concern is that AI systems, especially those driven by large datasets and personalization algorithms, can exploit cognitive biases and decision-making vulnerabilities in ways that are hidden or coercive.

The objective is to prevent the use of AI to covertly manipulate people into making choices they would not otherwise make, particularly in situations involving commercial, political, or social influence. These systems can undermine trust and autonomy by shaping preferences and behavior without users' knowledge or understanding. Article 5(1)(a) is not aimed at general persuasion or marketing but focuses specifically on harmful techniques that go beyond acceptable influence and impair the user's ability to act freely and rationally.

Main Concepts and Components

Three elements must be cumulatively met for the prohibition to apply: (1) the use of subliminal, purposefully manipulative, or deceptive techniques; (2) material distortion of behavior, impairing a person's ability to make an informed decision; and (3) significant harm caused or likely to be caused by that distortion.

"Subliminal" refers to stimuli presented below the threshold of conscious awareness, such as flashing images or inaudible audio messages. "Purposefully manipulative" techniques include those designed to exploit psychological vulnerabilities—for example, adaptive algorithms that escalate engagement through negative emotional triggers. "Deceptive" techniques mislead users about what the system does or the nature of the information it presents.

"Material distortion" implies that the user's ability to decide freely is appreciably undermined. The harm, whether physical, psychological, or financial, must be significant and not speculative. The prohibition applies regardless of whether the harm is intentional; what matters is the system's effects and the link between the technique and the outcome.

Interplay with Other EU Legal Acts

Article 5(1)(a) complements several areas of EU law. Under the **General Data Protection Regulation (GDPR)**, particularly Articles 5 and 6, systems that covertly profile users or process personal data without a legal basis may already be unlawful. The AI Act reinforces GDPR by addressing manipulation not just through data misuse but through behavioural impact.

It also intersects with the **Unfair Commercial Practices Directive (UCPD)**, which prohibits misleading or aggressive commercial conduct. While UCPD focuses on business-to-consumer practices, Article 5(1)(a) expands this to cover manipulative AI systems regardless of the sector, provided significant harm is reasonably likely.

The AI Act further aligns with the **Digital Services Act (DSA)**, especially regarding transparency and the prohibition of dark patterns in online platforms. AI systems that use manipulative design patterns—such as deceptive interface flows—may violate both DSA and the AI Act.

Together, these frameworks provide layered protection, with the AI Act uniquely targeting systemic and covert AI-driven manipulation at the technical level.

2. Exploitation of Vulnerable Groups (Art. 5.1.b)

Rationale and Objectives

Article 5(1)(b) of the AI Act aims to prevent the exploitation of individuals whose ability to make autonomous decisions is compromised due to vulnerabilities stemming from age, disability, or specific socio-economic conditions. The rationale is rooted in safeguarding human dignity, a foundational EU value under Article 1 of the Charter of Fundamental Rights. The provision targets AI systems that take advantage of these vulnerabilities in ways that distort decision-making and are likely to result in significant harm.

The objective is to protect groups with reduced resilience or limited capacity to recognize and resist manipulative AI influence. For example, children may be particularly susceptible to AI-powered gamification features that pressure them into making impulsive in-app purchases. Similarly, elderly or disabled persons could be misled by AI systems relying on voice or visual cues they cannot fully perceive or interpret. The AI Act responds to these risks by pre-emptively banning such exploitative systems, regardless of the intent of the provider or deployer.

Main Concepts and Components

The prohibition hinges on three cumulative criteria. First, the AI system must exploit vulnerabilities specifically linked to **age, disability, or socio-economic status**. These are defined narrowly to focus on users with impaired autonomy or understanding. Second, the system must have the **objective or effect of materially distorting behaviour**, meaning it significantly interferes with the user's ability to make informed choices. This is not about mere persuasion but about impairing agency in a substantial way. Third, the distortion must **cause or be reasonably likely to cause significant harm**, which may be psychological, financial, or physical.

Importantly, intent is not required—systems that unintentionally result in exploitative outcomes due to poor design or deployment still fall within the prohibition if they meet the conditions. However, lawful persuasion, marketing, or accessibility features do not fall under this rule unless they cross into exploitation and significant harm.

Interplay with Other EU Legal Acts

Article 5(1)(b) complements and reinforces several existing EU legal frameworks. It aligns closely with the **General Data Protection Regulation (GDPR)**, especially its emphasis on special protections for children's data and data minimization. Profiling that disproportionately affects vulnerable individuals without adequate safeguards may also contravene GDPR principles.

Moreover, the provision dovetails with the **Unfair Commercial Practices Directive (UCPD)**, which protects consumers—including vulnerable ones—from misleading or aggressive business practices. Article 5(1)(b) adds an AI-specific layer by explicitly targeting algorithmic systems designed (or functioning) to exploit such groups at scale.

Additionally, it supports the aims of the **European Pillar of Social Rights**, which calls for the protection and inclusion of persons with disabilities and those in precarious situations. The AI Act thus strengthens legal protection where other frameworks might not capture the technological nuance of AI-driven influence.

3. Social Scoring (Art. 5.1.c)

Rationale and Objectives

The prohibition on AI-enabled social scoring in Article 5(1)(c) reflects deep concerns over fairness, discrimination, and the abuse of data-based profiling. Inspired by systems seen in authoritarian contexts, where individuals are scored based on behaviors and traits, the EU aims to prevent such mechanisms from being used to marginalize, surveil, or punish citizens. These practices threaten fundamental EU values, including human dignity, equal treatment, and individual autonomy.

The primary objective is to ban AI systems that generate scores from behavioral or personality-based data and apply those scores in unrelated or disproportionate contexts—e.g., denying housing based on credit card behavior or restricting services based on political activity. The concern is not merely the use of data but the *cross-contextual* and *disproportionate* application of AI judgments that reinforce social stratification or arbitrary treatment.

Main Concepts and Components

Three cumulative elements define prohibited social scoring practices. First, the system must evaluate or classify individuals or groups over time based on social behavior or inferred personality traits. This includes tracking compliance, habits, lifestyle, or other subjective behavioral indicators.

Second, the score must result in either (i) *detrimental or unfavourable treatment in unrelated social contexts* or (ii) *unjustified or disproportionate treatment*. For instance, using shopping behavior to restrict access to education would be unrelated; barring someone from public transport for a minor administrative infraction would be disproportionate.

Third, the AI system must be “placed on the market,” “put into service,” or “used,” meaning the prohibition applies to both developers and deployers, whether public or private actors. Notably, intent is not a requirement; the actual or likely outcomes of the system determine its legality.

The prohibition does not apply to scoring that is *contextually relevant and proportionate*. For instance, evaluating driving history for insurance premiums, or customer reliability for e-commerce refunds, remains lawful where justified and not discriminatory.

Interplay with Other EU Legal Acts

Article 5(1)(c) complements a range of EU legal instruments. Under the **GDPR**, it supports principles of data minimization, fairness, and transparency. Profiling individuals across contexts for incompatible purposes can violate the purpose limitation principle and the restrictions on automated decision-making under Article 22 GDPR.

The provision also interacts with **anti-discrimination law**. If social scoring leads to direct or indirect discrimination based on protected characteristics—such as age, ethnicity, or political opinion—it may breach EU equality directives and Charter rights.

From a consumer law perspective, scoring systems by private actors that result in exploitative or misleading differentiation may fall under the **Unfair Commercial Practices Directive (UCPD)**. Particularly in business-to-consumer contexts, practices that materially distort consumer behavior or unfairly restrict access to services could trigger both AI Act and UCPD enforcement.

Finally, where social scoring intersects with **creditworthiness**, **employment**, or **migration** decisions, sectoral EU laws (like the Consumer Credit Directive or migration regulation) impose stricter conditions on data use and fairness. The AI Act's prohibition ensures such scoring systems do not circumvent these legal protections through opaque AI-based profiling.

4. Criminal risk assessment (Art. 5.1.d)

Rationale and Objectives

Article 5(1)(d) of the AI Act prohibits the use of AI systems to assess or predict whether a person will commit a criminal offence when such assessments are made **solely** based on profiling or on personality traits and characteristics. The underlying rationale is to uphold core EU values such as the **presumption of innocence**, **non-discrimination**, and **individual autonomy**. The use of AI to predict criminal behavior based on who a person is—rather than what they have done—poses serious risks of bias, stigmatization, and unjust restriction of fundamental rights.

The aim is to prevent preemptive criminalization, which may disproportionately affect marginalized communities and relies on speculative, often unverifiable indicators. By banning AI systems that make such risk assessments without any direct link to actual behavior or evidence, the AI Act seeks to limit overreach by both public and private actors and ensure that individuals are judged on their **actual actions**, not predictive traits.

Main Concepts and Components

The prohibition applies when three cumulative criteria are met. First, the AI system is used to **assess or predict the risk** of a person committing a crime. This includes determining likelihoods of future offenses or profiling for preemptive security actions. Second, the prediction is based **solely** on either profiling or assessing personality traits. "Profiling" refers to analyzing aspects of a person's identity or behavior to make predictions, while "personality traits and characteristics" covers features such as impulsivity or aggressiveness that may be inferred by AI models.

Third, the system must be **placed on the market**, **put into service**, or **used**—which broadens the scope to include developers, sellers, and users. Notably, this prohibition applies to both public entities (like law enforcement) and private actors (like insurers or hiring platforms) when these conditions are fulfilled.

However, AI systems used to support human assessments **based on objective and verifiable facts directly linked to criminal activity** are excluded from this prohibition. These systems may still be classified as "high-risk" under the AI Act and subject to specific obligations but are not outright banned.

Interplay with Other EU Legal Acts

The provision aligns closely with the **Law Enforcement Directive (LED)** and the **General Data Protection Regulation (GDPR)**. Under the LED, profiling that results in discrimination is already restricted, and Article 11(3) directly prohibits profiling leading to discriminatory treatment. The AI Act reinforces this by explicitly banning AI systems that rely exclusively on such profiling for crime prediction.

The prohibition also complements **Directive (EU) 2016/343** on the **presumption of innocence**, which safeguards individuals from being treated as guilty before proven so by law. While that directive applies from the moment of suspicion or accusation, the AI Act's scope is even broader, targeting prediction and prevention *before* any formal procedure begins.

Moreover, national **criminal procedure** and **data protection laws** continue to apply. Even where AI systems fall outside the Article 5(1)(d) prohibition, they may still be restricted by existing legal frameworks on due process, data minimization, or bias mitigation.

5. Untargeted facial image scraping (Art. 5.1.e)

Rationale and Objectives

Article 5(1)(e) of the AI Act prohibits the use of AI systems to create or expand facial recognition databases through the untargeted scraping of facial images from the internet or CCTV footage. The primary rationale is to safeguard individuals' rights to **privacy**, **data protection**, and **anonymity**, all of which are enshrined in the EU Charter of Fundamental Rights. The provision responds to growing concerns about the mass harvesting of biometric data without consent, often done covertly, and used to build databases for surveillance, profiling, or commercial purposes.

This practice creates what Recital 43 refers to as a "feeling of mass surveillance," infringing on the right to live freely without constant monitoring. The AI Act recognizes that untargeted scraping—especially from social media, video streams, or public camera feeds—can result in **gross violations of fundamental rights**, especially when individuals are unaware their images are being collected and processed.

Main Concepts and Components

The prohibition under Article 5(1)(e) hinges on four cumulative conditions. First, it applies to the **placing on the market**, **putting into service**, or **use** of an AI system. This ensures that both creators and users of such technologies are covered. Second, the system must have the **purpose of creating or expanding a facial recognition database**. It is not a general ban on facial recognition but focuses specifically on database generation.

Third, the image collection must occur through **untargeted scraping**—meaning the data is gathered indiscriminately, without specific targeting, context, or legal basis. Finally, the image sources must be either the **internet** (e.g., websites, social media) or **CCTV footage**. All four conditions must be met simultaneously for the prohibition to apply.

This prohibition is distinct from lawful and targeted facial image collection, such as acquiring consent-based images for access control or medical uses. It is narrowly focused on large-scale, indiscriminate harvesting operations that violate personal data rights. Tools that scrape billions of images to train facial recognition algorithms without users' knowledge are prime examples.

Interplay with Other EU Legal Acts

This prohibition aligns closely with the **General Data Protection Regulation (GDPR)**. Under the GDPR, facial images constitute **biometric data**, a special category of personal data requiring explicit consent for processing in most cases. Untargeted scraping likely violates GDPR principles of **lawfulness**, **transparency**, and **purpose limitation**, especially where individuals are unaware of or unable to object to data processing.

Additionally, this provision interacts with the **Law Enforcement Directive (LED)** when public authorities or their agents use such databases for surveillance. Even if law enforcement is not directly addressed under Article 5(1)(e), the same data may be subject to LED safeguards if accessed later for investigative purposes.

The rule also supports broader data minimization and fairness principles in **EU consumer protection** and **platform regulation** laws. For example, the **Digital Services Act (DSA)** promotes transparency and accountability in algorithmic data processing and reinforces protections against opaque and invasive data practices on online platforms.

Together, these frameworks ensure that AI systems relying on biometric data operate in a way that respects individual rights, prevents abuse, and protects against covert surveillance through mass data extraction.

6. Emotion recognition in workplace / education (Art. 5.1.f)

Rationale and Objectives

Article 5(1)(f) of the AI Act prohibits the use of AI systems designed to recognize emotions of individuals in **workplaces** and **educational institutions**. The rationale behind this provision is grounded in safeguarding the rights to **privacy**, **human dignity**, and **psychological integrity**. Emotion recognition technologies, which claim to detect internal states such as frustration, joy, anxiety, or deception based on facial expressions, tone of voice, or physiological signals, remain **scientifically contested** and **ethically sensitive**.

The core objective is to prevent misuse of such systems in environments where power asymmetries exist, and individuals may be compelled—implicitly or explicitly—to submit to emotional surveillance. In schools or workplaces, such monitoring could affect grading, hiring, promotions, or discipline. These environments are especially vulnerable to covert behavioral control or discrimination, with long-term effects on autonomy and mental well-being. The regulation thus seeks to preempt invasive uses of AI where individuals cannot freely consent or opt out without consequence.

Main Concepts and Components

This prohibition applies when AI systems are (1) used to **infer emotional states**, (2) within **educational or work-related settings**, and (3) not exempt under limited, well-defined **medical or safety exceptions**.

Emotion recognition refers to systems that analyze behavioral, physiological, or biometric data to deduce emotional states. While commonly marketed in recruitment, training, or learning optimization, these systems can be unreliable and culturally biased. For example, using facial expressions to measure student engagement or employee honesty is both scientifically weak and ethically questionable.

The scope is limited to uses **in educational institutions or at work**, whether by public or private actors. It covers systems used on **employees, students, candidates, or trainees**, whether online or in-person.

The prohibition includes systems built into online platforms, proctoring tools, or office software if deployed within the scope of education or employment.

There are **explicit exceptions** where emotion recognition is used strictly for **medical or safety reasons**, such as detecting emotional distress in mental health treatment or fatigue monitoring in high-risk jobs (e.g., aviation, transport). Even in those cases, compliance with other applicable laws is mandatory.

Interplay with Other EU Legal Acts

The AI Act's prohibition interacts closely with the **General Data Protection Regulation (GDPR)**, particularly as emotion data may constitute **biometric or inferred personal data**. Under GDPR, such processing often requires explicit consent and a clear legal basis, which are hard to establish in dependent relationships like employment or schooling. This reinforces the AI Act's intent to prevent coercive or opaque use of sensitive AI in power-imbalanced settings.

The rule also supports the **Charter of Fundamental Rights**, especially Articles 1 (dignity), 8 (data protection), and 21 (non-discrimination), ensuring individuals are not judged by unverifiable emotional metrics. Moreover, **occupational health and safety** directives and **education law** in Member States may add further restrictions or safeguards—even in cases exempt from the AI Act.

In sum, this provision ensures that AI stays out of intimate, human-centered processes like teaching, learning, and working—unless strictly justified, safe, and rights-compliant.

7. Biometric categorisation (Art. 5.1.g)

Rationale and Objectives

Article 5(1)(g) of the AI Act prohibits AI systems that categorize individuals based on their biometric data in order to deduce or infer certain **"sensitive" characteristics**, such as race, political opinion, religious beliefs, sexual orientation, and more. The rationale behind this prohibition is to prevent discriminatory profiling, the erosion of personal dignity, and the use of AI in ways that can lead to exclusion, stigmatization, or targeting based on deeply personal traits.

The AI Act recognizes that biometric data—such as facial geometry, iris patterns, or gait—can be processed to infer far more than mere identity. These systems can expose or predict protected attributes without individuals' consent or awareness, potentially leading to unfair treatment. This raises profound concerns about privacy, equality, and non-discrimination. The objective is to **ban categorisation AI systems** that turn biometric data into a tool for sorting people into social, political, or ideological boxes.

Main Concepts and Components

Three key elements define this prohibition. First, it applies to **biometric categorisation systems**, meaning AI systems that process biometric data to assign individuals into specific groups or categories. These categories do not have to be formally labelled; the act of sorting based on inferred traits is enough.

Second, the categorisation must be **individualized**—targeting natural persons rather than groups or statistical aggregates. It's not about classifying a crowd's demographics for statistical analysis, but rather assigning sensitive attributes to specific people based on biometric input.

Third, the inference must relate to **"special categories"** of attributes: race or ethnic origin, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation. These are aligned with Article 9(1) GDPR, which already defines such data as requiring heightened protection.

Notably, the prohibition does not apply to **non-sensitive biometric classification**, such as categorizing by age group or eye color, provided the data was acquired lawfully and used within regulated contexts (e.g. medical or forensic applications). It also excludes anonymized or statistical categorisation, as long as no individual can be identified or inferred.

Interplay with Other EU Legal Acts

This provision reinforces and expands on protections established in the **General Data Protection Regulation (GDPR)**. Article 9(1) GDPR restricts the processing of sensitive data unless strict conditions are met. Article 5(1)(g) of the AI Act effectively closes a loophole where sensitive attributes could be inferred from biometric data without explicitly being collected as such.

It also interacts with the **Law Enforcement Directive (LED)**, particularly Article 11(3), which prohibits profiling that leads to discrimination based on sensitive personal data. The AI Act functions as lex specialis in this context, imposing a categorical prohibition on certain AI systems, not just conditions or safeguards.

In addition, the rule supports **EU equality law** and the **Charter of Fundamental Rights**, notably Article 21 (non-discrimination). The potential misuse of biometric AI to sort people by race or beliefs runs counter to these core principles, especially in contexts like access to services, hiring, or surveillance.

8. Real-Time Remote Biometric Identification in Public Spaces (Art.5.1.h)

Rationale and Objectives

Article 5(1)(h) of the AI Act prohibits the use of **real-time remote biometric identification systems** in **publicly accessible spaces** for **law enforcement purposes**, with narrow exceptions. The rationale is rooted in the protection of **privacy**, **freedom of movement**, **freedom of assembly**, and **non-discrimination**, which are at risk when biometric surveillance tools are used without strict controls. These technologies, by identifying individuals at a distance using facial or gait recognition, enable mass tracking and risk creating a culture of constant surveillance.

Recital 32 of the AI Act emphasizes that the intrusive nature of real-time RBI can chill civil liberties and disproportionately affect vulnerable groups—particularly when inaccuracies or biases in the technology result in **false positives**, especially for racial minorities or persons with disabilities. The regulation thus seeks to prevent disproportionate, unaccountable surveillance while still allowing limited, justified use in circumstances where **public safety outweighs the risk to fundamental rights**.

Main Concepts and Components

The prohibition applies when four criteria are met cumulatively: (1) the system is an **RBI system**; (2) it operates in **real-time**; (3) in a **publicly accessible space**; and (4) it is used for **law enforcement purposes**.

“Remote biometric identification” refers to the automatic identification of individuals without their active participation, using facial recognition or similar technologies. “Real-time” implies the system functions live, matching biometric input against a reference database instantaneously or with minimal delay. “Publicly accessible spaces” include streets, parks, stadiums, or transit hubs. The term “for law enforcement purposes” extends to both police and entities acting on their behalf, such as transport companies under law enforcement instruction.

Three exceptions are permitted under Article 5(1)(h)(i–iii):

1. **Targeted searches for victims** of abduction, human trafficking, or sexual exploitation, and missing persons.

2. **Prevention of specific, imminent threats** to life or physical safety, including terrorism.
3. **Identification of suspects** in connection with serious crimes listed in Annex II of the AI Act, punishable by four years or more.

These exceptions are tightly defined and require **national legislation**, prior **judicial or independent authority approval**, and a **Fundamental Rights Impact Assessment (FRIA)**. Use must be targeted—not blanket scanning of the public—and remain proportionate, geographically and temporally limited.

Interplay with Other EU Legal Acts

This prohibition intersects with the **Law Enforcement Directive (LED)** and **GDPR**, which govern the processing of biometric data and mandate respect for individual rights. The AI Act acts as **lex specialis** in this context, setting stricter, purpose-specific rules for RBI systems used by or on behalf of law enforcement.

Moreover, any authorized use must still comply with **data protection principles** such as necessity, proportionality, and purpose limitation. Even lawful RBI under Article 5(1)(h) must meet all procedural and substantive conditions laid out in both the AI Act and national law.

Enforcement and Consequences

Violating these prohibitions carries **severe penalties** under the AI Act. The maximum fine for using a prohibited AI practice can reach **€35 million or 7% of global annual turnover**, whichever is higher. These are among the highest sanctions available under EU law, signaling the seriousness with which the EU treats these banned applications. Supervisory authorities in each member state will be empowered to investigate and sanction non-compliance, including by banning systems from entering or remaining on the EU market.

Figure 1: EU AI Act Timeline

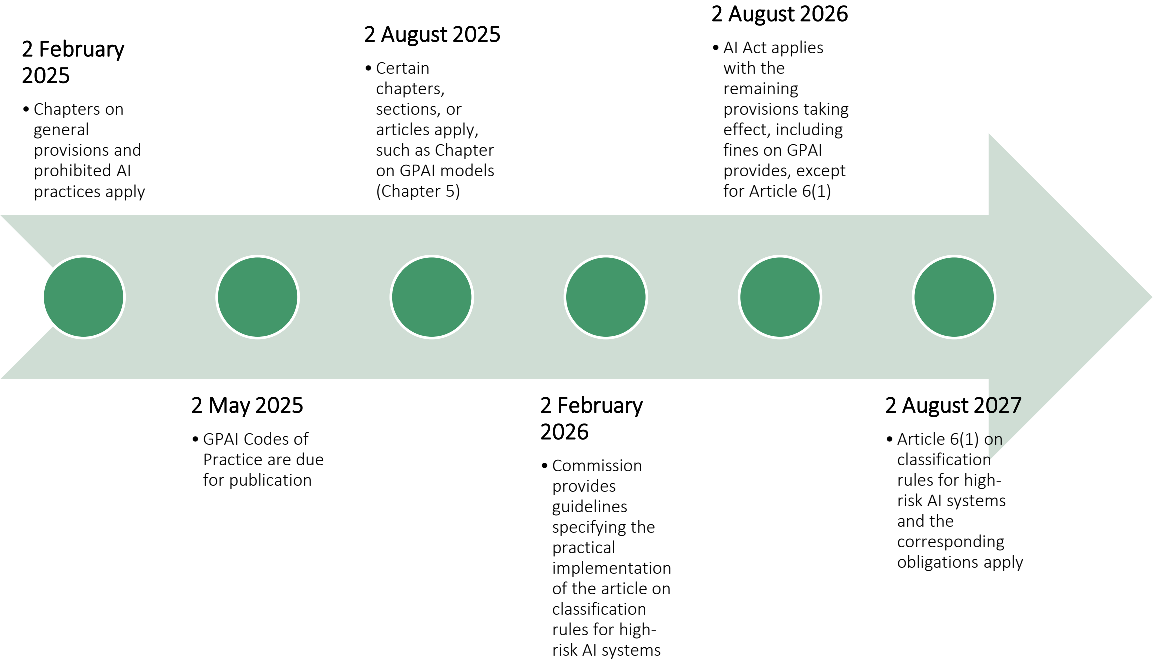
2 February 2025
• Chapters on general provisions and prohibited AI practices apply
2 August 2025
• Certain chapters, sections, or articles apply, such as Chapter on GPAI models (Chapter 5)
2 August 2026
• AI Act applies with the remaining provisions taking effect, including fines on GPAI provides, except for Article 6(1)
2 May 2025
• GPAI Codes of Practice are due for publication
2 February 2026
• Commission provides guidelines specifying the practical implementation of the article on classification rules for high-risk AI systems
2 August 2027
• Article 6(1) on classification rules for high-risk AI systems and the corresponding obligations apply

The EU AI Act introduces a comprehensive penalty framework aimed at ensuring compliance with its requirements. The fines under the Act are structured to be dissuasive and proportionate, with maximum thresholds defined both as fixed monetary amounts and as percentages of the global annual turnover of the offending undertaking. The severity of penalties varies based on the type and gravity of the infringement.

Figure 2: EU AI Act Penalty Structure

VIOLATION TYPE	MAXIMUM FINE
PROHIBITED AI PRACTICES (ART. 5)	€35 million or 7% of worldwide annual turnover
BREACHES OF KEY OBLIGATIONS (E.G., FOR PROVIDERS, IMPORTERS, DEPLOYERS)	€15 million or 3% of worldwide annual turnover
MISLEADING OR INCOMPLETE INFORMATION PROVIDED TO REGULATORS	€7.5 million or 1% of worldwide annual turnover
VIOLATIONS BY PROVIDERS OF GENERAL-PURPOSE AI (GPAI) MODELS	€15 million or 3% of worldwide annual turnover
EU INSTITUTIONS OR BODIES BREACHING APPLICABLE PROVISIONS	Administrative fines imposed by the EDPS (variable)

Prohibited AI Practices (Article 5 Violations):
The most severe category under the Act, breaches here attract fines up to €35 million or 7% of the worldwide annual turnover. This includes engaging in AI uses banned outright by the Act, such as manipulative AI or social scoring systems.

Obligations on Operators and Notified Bodies:
Non-compliance with operational responsibilities—such as risk management, data governance, and human oversight (under Articles 16, 21–24, 26, 31–34, and 50)—can result in fines up to €15 million or 3% of global turnover.

Providing Misleading or Incomplete Information (Art. 91):
Entities that provide false, incomplete, or misleading information to the Commission or other regulators may face fines up to €7.5 million or 1% of their worldwide turnover.

GPAI Model Providers (Art. 101):
These providers are subject to direct Commission enforcement. Fines of up to €15 million or 3% of global turnover apply for violations such as failing to provide documentation, comply with access requests, or respond to systemic risk assessments.

EU Institutions and Agencies:
For institutions within the EU, the European Data Protection Supervisor (EDPS) is empowered under Article 100 to impose administrative fines based on considerations like the severity of breach and mitigation efforts.

Special Provisions for SMEs

The Act explicitly aims to safeguard small and medium-sized enterprises (SMEs) by capping their fines at the lower of the percentage or absolute amount thresholds. This tiered approach balances enforcement with support for innovation and economic viability among smaller actors.

The list of prohibited practices is deliberately narrow, but its impact is far-reaching. For developers, it means certain types of AI—especially those involving biometric surveillance, psychological profiling, or behavior manipulation—are off-limits in the EU regardless of technical performance or commercial potential.

For governments and law enforcement, the Act sends a clear signal: AI must respect the principles of necessity and proportionality, and must never cross into blanket surveillance or arbitrary discrimination.

And for the public, the list functions as a rights guarantee—an assurance that some lines cannot be crossed, no matter how powerful or profitable a technology may be.

Multi-Level Enforcement Structure and Authority Roles

The EU AI Act introduces a layered enforcement regime involving EU-wide bodies and national authorities. At the top sits the European AI Office, which oversees GPAI model compliance, coordinates between Member States, and ensures regulatory consistency. National competent authorities are responsible for enforcement related to high-risk AI systems, including conducting investigations and issuing penalties. The European Data Protection Supervisor (EDPS) holds jurisdiction over EU institutions' use of AI. This multi-tiered approach mirrors other EU frameworks like GDPR, but coordination risks remain, especially where enforcement overlaps or diverging interpretations among Member States could create fragmented compliance expectations for global businesses.

Future Regulatory Trends

Given the EU AI Act's evolving nature, future regulatory trends are likely to emphasize interpretative clarity and sector-specific guidance. As implementation begins, expect an increase in delegated acts and harmonized standards, particularly for GPAI models and high-risk systems. Regulatory sandboxes and innovation hubs may proliferate to support compliance experimentation without punitive risk. Trends may also include expanding the Act's extraterritorial reach, enhanced transparency requirements, and integrated governance with digital, privacy, and product safety laws. Businesses should prepare to adapt to a living framework—one that will likely evolve with new technologies, litigation outcomes, and growing cross-border enforcement coordination.

Conclusion

The EU AI Act does not aim to regulate all AI uniformly, but it does establish **clear red lines**. The list of prohibited AI practices reflects the EU's commitment to values-driven governance: certain technologies, no matter how sophisticated, must be rejected if they undermine the rights, dignity, or safety of individuals.

By banning practices like real-time biometric surveillance, manipulative AI, and social scoring, the EU sets a global precedent—showing that ethical boundaries can be drawn even in the face of rapid technological change. These bans are not just legal restrictions; they are statements of principle about the kind of digital future Europe is willing to accept.

Glossary

Act or EU AI Act: European Union Artificial Intelligence Act

AI: Artificial Intelligence

Board: European Union Artificial Intelligence Board

EU: European Union

SME: Small and Medium-Sized Enterprise

AI & Partners
Amsterdam - London - Singapore

How can we help?

AI & Partners – 'AI That You Can Trust'

At AI & Partners, we're here to help you navigate the complexities of the EU AI Act, so you can focus on what matters—using AI to grow your business. We specialize in guiding companies through compliance with tailored solutions that fit your needs. Why us? Because we combine deep AI expertise with practical, actionable strategies to ensure you stay compliant and responsible, without losing sight of your goals. With our support, you get AI you can trust—safe, accountable, and aligned with the law.

To find out how we can help you, email contact@ai-and-partners.com or visit https://www.ai-and-partners.com.