

EU AI Act: Risk-Based Regulatory Framework

A Guide to AI Restrictions, Requirements, and Responsibilities

What restrictions, requirements, and responsibilities should you understand to develop, deploy, and use AI tools in your business operations?

Overview: The EU AI Act

As an employee working with AI, the **EU AI Act** [tn500] (Regulation (EU) 2024/1689) [1] is the legal framework you need to understand. It is the world's first comprehensive regulation for **artificial intelligence** [tn501] and it **governs** [te009] every AI system your organization develops, deploys, or uses — through a risk-based approach that categorizes AI into tiers based on its potential impact on fundamental rights and safety. At the highest tier, certain AI practices are **classified above** [te001] all others as outright **prohibited** [tn100] (Art. 5) — the Act **prohibits** [te007] eight categories of unacceptable-risk practices that are banned, with narrow exceptions for law enforcement. Systems your organization classifies as **high-risk** [tn200] (Art. 6, Annex III) **must comply with** [te002] stringent **compliance requirements** [tn300] (Art. 8–15) before being placed on the market. Each high-risk domain **falls under** [te005] specific obligations that **require** [te006] conformity assessment. The regulation is **enforced from** [te003] a **phased enforcement timeline** [tn400] — the earliest prohibitions and AI literacy requirements already apply since February 2025, **followed by** [te008] successive phases, with full enforcement by August 2027. The Act **applies to** [te004] **providers** [tn502] (organizations developing AI), **deployers** [tn503] (organizations

using AI), and any entity whose AI system outputs affect persons in the EU. Understanding your role in this framework **enables** [te010] you to operate AI responsibly and in compliance across all Member States.

Prohibited Practices [tn100]: Art. 5 — Unacceptable Risk

Article 5 of the AI Act identifies AI practices at a **risk level above** [te101] all other categories, posing an unacceptable risk to fundamental rights. These practices are banned outright, with narrow exceptions for law enforcement under strict judicial safeguards [1].

- **Subliminal Manipulation** [tn101]: 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 behavior, thereby causing or being reasonably likely to cause significant harm
- **Vulnerability Exploitation** [tn102]: AI systems that exploit vulnerabilities due to age, disability, or a specific social or economic situation to materially distort behavior, causing or likely to cause significant harm
- **Social Scoring** [tn103]: AI systems used to evaluate or classify natural persons or groups based on their social behavior or known, inferred, or predicted personal characteristics over time, where the resulting social score leads to detrimental or unfavorable treatment in social contexts unrelated to those in which the data was originally collected, or treatment disproportionate to the gravity of their social behavior
- **Predictive Policing** [tn104]: AI systems making risk assessments of individual natural persons to predict criminal offenses based solely on profiling or personality traits, unless supporting human assessment grounded in objective, verifiable facts directly linked to criminal activity
- **Untargeted Facial Scraping** [tn105]: AI systems that build or expand facial recognition databases through untargeted scraping of facial

images from the internet or CCTV footage

- **Workplace and Education Emotion Inference** [tn106]: AI systems that infer the emotions of natural persons in workplace and educational settings, except where intended for medical or safety purposes
- **Biometric Categorization** [tn107]: AI systems that categorize natural persons individually based on biometric data to deduce or infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation, except for lawful labeling or filtering of biometric datasets or law enforcement categorization
- **Real-Time Remote Biometric Identification** [tn108]: Use of real-time remote biometric identification systems in publicly accessible spaces for law enforcement purposes, except under narrow exceptions such as targeted search for missing persons, prevention of imminent terrorist threats, or locating suspects of serious criminal offenses, each subject to prior judicial authorization

High-Risk Classification [tn200]: Art. 6 and Annex III Domains

Article 6 classifies AI systems as high-risk via two pathways [1]. First, AI systems that are safety components of products covered by Union harmonisation legislation listed in Annex I (such as medical devices, machinery, toys, lifts, and vehicles) are high-risk when the product requires third-party conformity assessment (Art. 6(1)). Second, standalone AI systems used in the eight sensitive domains listed in Annex III below are high-risk (Art. 6(2)). In both cases, systems that pose a significant risk of harm to health, safety, or fundamental rights **must comply with** [te102] the full set of **compliance requirements** [tn300] prior to **placing on the market** [tn504] or putting into service.

- **Biometrics** [tn201]: Remote biometric identification (non-real-time), emotion recognition systems, and biometric categorization of natural persons

- **Critical Infrastructure** [tn202]: AI used as safety components in management and operation of roads, water supply, gas, heating, and electricity infrastructure
- **Education and Vocational Training** [tn203]: AI systems determining access to or admission into educational institutions, assessing learning outcomes of students, monitoring examination conduct, and evaluating appropriate education levels for individuals
- **Employment and Worker Management** [tn204]: AI systems used in recruitment, candidate selection, employment-related decisions, task allocation, performance and behavior monitoring of workers
- **Essential Private and Public Services** [tn205]: AI systems evaluating creditworthiness, credit scoring, risk assessment for life and health insurance pricing, and prioritization of emergency services dispatch
- **Law Enforcement** [tn206]: AI systems for individual risk assessments of natural persons, polygraph-like tools and lie detection, evaluation of evidence reliability, and crime analytics including profiling for prediction and detection
- **Migration, Asylum and Border Control** [tn207]: AI systems for lie detection in asylum and visa procedures, risk assessment of irregular migration threats, and screening and security risk evaluation of natural persons
- **Administration of Justice and Democracy** [tn208]: AI systems intended to assist judicial authorities in researching and interpreting facts and the law and in applying the law, as well as systems used to influence the outcome of elections or referendums or the voting behavior of natural persons
- **High-Risk Derogation** [tn209] (**Art. 6(3)**): An AI system listed in Annex III is NOT considered high-risk if it does not pose a significant risk to health, safety, or fundamental rights, provided it meets at least one condition: (a) it performs a narrow procedural task, (b) it improves the result of a previously completed human activity, (c) it detects decision-making patterns without replacing human assessment, or (d) it performs a preparatory task for an Annex III assessment. However, an Annex III system that performs profiling of natural persons is always considered high-risk regardless of these conditions. Providers relying on this

derogation must document their assessment and register in the EU database (Art. 49)

Compliance Requirements [tn300] for High-Risk AI Systems: Art. 8–15

Providers [tn502] of **high-risk** [tn200] AI systems must comply with eight core requirements before **placing on the market** [tn504], covering the entire lifecycle from design through deployment and post-market monitoring [1].

- **Risk Management System** [tn301] (**Art. 9**): A continuous, iterative process planned and maintained throughout the entire lifecycle of the high-risk AI system, requiring systematic identification, analysis, estimation, and evaluation of risks and adoption of appropriate and targeted risk management measures. The risk management system **informs** [te301] the data governance practices that follow
- **Data and Data Governance** [tn302] (**Art. 10**): Training, validation, and testing data sets must meet quality criteria including relevance, representativeness, completeness, and statistical properties appropriate to the intended purpose, with governance practices covering data collection, preparation, and bias examination. These practices are **documented in** [te302] the technical documentation
- **Technical Documentation** [tn303] (**Art. 11**): Comprehensive technical documentation drawn up before the system is placed on the market or put into service, kept up-to-date, and made available to national competent authorities upon request to demonstrate compliance
- **Record-Keeping** [tn304] (**Art. 12**): High-risk AI systems must be designed to allow automatic recording of events (logs) over their lifetime, enabling traceability of the system's functioning and facilitating post-market monitoring

- **Transparency and Provision of Information** [tn305] (**Art. 13**): High-risk AI systems must be designed to operate with sufficient transparency for deployers to interpret the system's output and use it appropriately, accompanied by clear and adequate instructions for use. Transparency **enables** [te303] effective human oversight
- **Human Oversight** [tn306] (**Art. 14**): High-risk AI systems must be designed with appropriate human-machine interface tools so that they can be effectively overseen by natural persons during use, including the ability to understand capacities and limitations, monitor operation, and intervene or interrupt the system
- **Accuracy, Robustness and Cybersecurity** [tn307] (**Art. 15**): High-risk AI systems must achieve appropriate levels of accuracy, robustness, and cybersecurity throughout their lifecycle, with resilience to errors, faults, inconsistencies, and adversarial attacks including data poisoning and model manipulation. These properties are **validated by** [te304] the overall conformity assessment
- **Overall Compliance** [tn308] (**Art. 8**): High-risk AI systems must comply with all requirements set out in Articles 9–15, taking into account the generally acknowledged state of the art, the intended purpose of the system, and relevant harmonised standards or common specifications, followed by a conformity assessment procedure before placing on the market

Limited Risk Obligations [tn550]: Art. 50 — Transparency Requirements

Article 50 establishes transparency obligations for AI systems that interact directly with natural persons or generate synthetic content. These **apply to** [te550] specific categories of AI systems — regardless of their risk classification — that pose risks of deception or manipulation. As an employee, understanding these obligations is critical when deploying chatbots, content generation tools, or biometric systems in customer-facing or internal operations [1].

- **Human Interaction Disclosure** [tn551]: Providers must ensure that AI systems intended to interact directly with natural persons (such as chatbots) are designed so that the persons concerned are informed they are interacting with an AI system, unless this is obvious from the circumstances and context of use. This **requires disclosure by** [te551] the provider at the point of first interaction
- **Synthetic Content Marking** [tn552]: Providers of AI systems generating synthetic audio, image, video, or text content must ensure that outputs are marked in a machine-readable format and detectable as artificially generated or manipulated. This includes general-purpose AI systems. The technical solutions must be effective, interoperable, robust, and reliable. This obligation does not apply to AI systems that perform an assistive function for standard editing or that do not substantially alter the input data provided by the deployer
- **Emotion Recognition and Biometric Disclosure** [tn553]: Deployers of emotion recognition systems or biometric categorization systems must inform the natural persons exposed to the system of its operation, and process personal data in accordance with GDPR and related data protection regulations. This obligation does not apply to AI systems used for law enforcement
- **Deep Fake Disclosure** [tn554]: Deployers of AI systems that generate or manipulate image, audio, or video content constituting a deep fake must disclose that the content has been artificially generated or manipulated. This **extends to** [te552] text published on matters of public interest, which must also be disclosed as AI-generated, unless the content has undergone human editorial review

General-Purpose AI Models [tn600]: Art. 51–56 — GPAI Model Obligations

Chapter V introduces a dedicated regulatory regime for general-purpose AI (GPAI) models — large-scale models trained on broad data that can serve a variety of downstream tasks, such as large language models. These rules **apply separately from** [te600] the risk-based classification for AI systems, recognizing that GPAI models **underpin** [te601] many downstream AI systems. As an employee, if your organization develops, deploys, or integrates foundation models, you must understand both the base obligations and the additional systemic risk requirements [1].

- **Systemic Risk Classification** [tn601] (**Art. 51**): A GPAI model is classified as presenting systemic risk if it has high-impact capabilities (evaluated via benchmarks) or if the Commission designates it based on criteria in Annex XIII. A model is presumed to have high-impact capabilities when the cumulative computation used for training exceeds 10^{25} floating-point operations. This classification **triggers additional obligations under** [te602] Article 55
- **Technical Documentation** [tn602] (**Art. 53**): GPAI model providers must draw up and maintain technical documentation including training and testing processes and evaluation results, and make this available to the AI Office and national competent authorities upon request
- **Downstream Provider Information** [tn603] (**Art. 53**): GPAI model providers must make available information and documentation to downstream providers integrating the model into AI systems, enabling them to understand capabilities, limitations, and their own compliance obligations. This information must contain, at minimum, the elements set out in Annex XII
- **Copyright Compliance** [tn604] (**Art. 53**): GPAI model providers must put in place a policy to comply with Union copyright law, including identifying and complying with reservations of rights expressed pursuant to the Copyright Directive (EU) 2019/790, and publish a sufficiently detailed summary of the training data content
- **Systemic Risk Obligations** [tn605] (**Art. 55**): Providers of GPAI models with systemic risk must additionally perform model evaluations including adversarial testing, assess and mitigate systemic risks at Union level, track and report serious incidents to the AI Office, and ensure adequate

cybersecurity protection for the model and its infrastructure. These obligations **go beyond** [te603] the base requirements of Article 53

- **Codes of Practice** [tn606] **(Art. 56)**: The AI Office facilitates codes of practice at Union level to demonstrate compliance with GPAI obligations. These codes cover technical documentation standards, training data summaries, systemic risk identification, and risk assessment procedures. Codes of practice must be ready by 2 May 2025, and compliance with approved codes **grants presumption of conformity with** [te604] Articles 53 and 55 until harmonised standards are published
- **Open-Source Exception** [tn607]: GPAI models released under free and open-source licences, whose parameters, architecture, and usage information are made publicly available, **are exempt from** [te906] the technical documentation and downstream information obligations (Art. 53(1)(a)-(b)). However, this exception does not apply to GPAI models with systemic risk — they must comply with all obligations regardless of licensing

Obligations by Role [tn700]: Art. 16–27 — Who Must Do What

The AI Act assigns distinct obligations to each actor in the AI value chain. As an employee, your responsibilities depend on whether your organization **acts as** [te700] a provider, deployer, importer, or distributor of high-risk AI systems. Understanding your role determines which compliance requirements apply to your daily work [1].

- **Provider Obligations** [tn701] **(Art. 16–21)**: Providers bear the heaviest obligations: ensure compliance with all requirements (Art. 8–15), implement a quality management system (Art. 17), maintain technical documentation for 10 years (Art. 18), keep automatically generated logs (Art. 19), take corrective actions for non-conforming systems (Art. 20), **cooperate with** [te701] national competent authorities (Art. 21), affix **CE**

- marking [tn509] (Art. 48), establish **post-market monitoring** [tn507] (Art. 72), and undergo conformity assessment before placing on the market
- **Deployer Obligations** [tn702] (**Art. 26**): Deployers must use AI systems in accordance with instructions for use, assign human oversight to competent persons, ensure input data is relevant and representative, monitor operation and report risks to providers, keep logs for at least six months, report any **serious incident** [tn508] to the market surveillance authority (Art. 73), inform workers' representatives before workplace deployment, inform affected persons that they are subject to AI decision-making, and **carry out** [te702] a fundamental rights impact assessment (Art. 27) when deploying as a public body or in essential services
 - **Importer Obligations** [tn703] (**Art. 23**): Importers must verify that the provider has carried out the conformity assessment, that technical documentation exists, that the CE marking is affixed, and that an authorised representative has been appointed. Importers must not place non-conforming systems on the market, must keep documentation for 10 years, and must cooperate with competent authorities
 - **Distributor Obligations** [tn704] (**Art. 24**): Distributors must verify CE marking, EU declaration of conformity, and instructions for use before making a system available. They must not make non-conforming systems available and must inform the provider or importer of any identified non-conformity or risk
 - **Value Chain Responsibilities** [tn705] (**Art. 25**): Any distributor, importer, deployer, or third party **becomes a provider** [te703] (and assumes all provider obligations) if they: put their name or trademark on a high-risk AI system, make a substantial modification, or modify the intended purpose so it becomes high-risk. This prevents compliance gaps through relabelling or modification
 - **Fundamental Rights Impact Assessment** [tn706] (**Art. 27**): Public-body deployers and private entities providing essential services must perform an assessment of the impact on fundamental rights before deploying high-risk AI systems. The assessment covers: description of processes, period and frequency of use, categories of affected persons, specific risks of harm, human oversight measures, and **governance mechanisms for** [te704] complaint handling. Results must be notified to the market surveillance authority

- **Authorised Representative** [tn707] (**Art. 22**): Providers of high-risk AI systems established in third countries (outside the EU) must appoint, by written mandate, an authorised representative established in the Union before making their systems available on the EU market. The authorised representative verifies that the EU declaration of conformity and technical documentation have been drawn up, keeps copies at the disposal of authorities for 10 years, and cooperates with the AI Office and competent authorities on behalf of the provider

Cross-Cutting Provisions [tn800]: Obligations and Voluntary Measures Across All AI

Certain obligations **apply across all risk tiers** [te800], regardless of whether an AI system is classified as prohibited, high-risk, limited-risk, or minimal-risk. These cross-cutting provisions — some mandatory, some voluntary — form the broader context that every employee working with AI must understand [1]. Note that the vast majority of AI systems in everyday business use (such as spam filters, inventory optimization, or recommendation systems) fall into the minimal-risk category and face no mandatory obligations beyond AI literacy — they may voluntarily adopt codes of conduct.

- **AI Literacy** [tn801] (**Art. 4**): Providers and deployers must ensure, to their best extent, a sufficient level of AI literacy of their staff and other persons dealing with the operation and use of AI systems on their behalf. This obligation **takes into account** [te801] the technical knowledge, experience, education, and training of the persons concerned, as well as the context in which the AI systems are to be used. AI literacy applies from 2 February 2025 — it is one of the earliest obligations to take effect
- **GDPR and Data Protection** [tn802]: The AI Act operates without prejudice to GDPR (Regulation 2016/679), LED (Directive 2016/680), and the ePrivacy Directive (2002/58/EC). Any AI system processing personal data

must comply with existing data protection law in addition to AI Act requirements. Data protection impact assessments under GDPR Article 35 complement the fundamental rights impact assessments under AI Act Article 27

- **Non-Discrimination** [tn803]: AI systems must be developed and used in compliance with Union non-discrimination law. Bias in training data, model design, or deployment can lead to discriminatory outcomes against protected groups based on race, gender, disability, age, or other characteristics. Providers of high-risk systems must specifically examine data sets for possible biases (Art. 10(2)(f)) and adopt measures to detect, prevent, and mitigate them
- **Voluntary Codes of Conduct** [tn804]: Providers of non-high-risk AI systems are encouraged to **voluntarily apply** [te802] the mandatory requirements for high-risk systems, including creating codes of conduct covering AI literacy, environmental sustainability, inclusive design, diversity in development teams, and accessibility. These voluntary commitments should be based on clear objectives and measurable indicators
- **Right to Explanation** [tn805]: Affected persons have the right to obtain an explanation when a deployer's decision is based mainly on the output of certain high-risk AI systems and that decision produces legal effects or similarly significantly affects them. The explanation must be clear and meaningful, **enabling affected persons to** [te803] exercise their rights. Any person may lodge a **complaint** [tn807] with a market surveillance authority (Art. 85). Persons reporting infringements are protected under **whistleblower protection** [tn806] (Directive 2019/1937). These rights apply in addition to existing rights under GDPR Article 22

Governance and Enforcement [tn900]: Institutional Framework and Penalties

The AI Act establishes a multi-layered governance framework at EU and national levels, with the **EU AI Office** [tn901] at its centre. Enforcement is carried out through market surveillance authorities in each Member State, with tiered administrative fines for non-compliance. As an employee, understanding the governance structure helps you know who **oversees** [te900] your organization's AI activities and what penalties apply [1].

- **EU AI Office** [tn901] (**Art. 64**): The AI Office, established within the European Commission, develops Union expertise and capabilities in AI. It has the powers of a market surveillance authority for GPAI model supervision, can investigate infringements, request documentation, conduct evaluations, and **enforce compliance for** [te901] general-purpose AI model providers at Union level
- **European AI Board** [tn902] (**Art. 65–66**): The European Artificial Intelligence Board, composed of one representative per Member State, advises the Commission and Member States on consistent and effective application of the Regulation. The Board coordinates among national authorities, collects technical expertise, provides guidance on GPAI enforcement, and **supports harmonisation of** [te902] administrative practices across the Union. It is supported by an **advisory forum** [tn909] (Art. 67) of industry, civil society, academia, and SME stakeholders
- **National Competent Authorities** [tn903] (**Art. 70**): Each Member State must designate at least one notifying authority and one market surveillance authority. These authorities exercise powers independently and impartially, with adequate AI expertise, data protection knowledge, and cybersecurity competence. Market surveillance authorities **carry out enforcement at** [te903] the national level, including inspections, document requests, and corrective actions. All high-risk AI systems must be registered in the **EU database** [tn908] (Art. 71) before being placed on the market
- **Penalty Tier 1: Prohibited Practices** [tn904] (**Art. 99(3)**): Non-compliance with the prohibition of AI practices referred to in Article 5 **penalises violations of** [te907] prohibited practices, subject to administrative fines of up to EUR 35 000 000 or, for undertakings, up to 7% of total worldwide annual turnover for the preceding financial year, whichever is higher. For

SMEs and start-ups, the fine is capped at the lower of the percentage or the fixed amount

- **Penalty Tier 2: Operator Obligations** ^[tn905] (**Art. 99(4)**): Non-compliance with operator obligations (provider Art. 16, deployer Art. 26, importer Art. 23, distributor Art. 24), notified body requirements, or transparency obligations (Art. 50) is subject to fines of up to EUR 15 000 000 or 3% of total worldwide annual turnover, whichever is higher. This tier **covers the majority of** ^[te904] day-to-day compliance failures
- **Penalty Tier 3: Incorrect Information** ^[tn906] (**Art. 99(5)**): Supplying incorrect, incomplete, or misleading information to notified bodies or national competent authorities is subject to fines of up to EUR 7 500 000 or 1% of total worldwide annual turnover, whichever is higher. This applies to requests during conformity assessments, market surveillance investigations, or any authority inquiry
- **GPAI Model Fines** ^[tn907] (**Art. 101**): The Commission may impose fines on GPAI model providers not exceeding 3% of annual worldwide turnover or EUR 15 000 000, whichever is higher, for infringement of GPAI obligations, failure to comply with information requests, or failure to provide access to the model for evaluation. The Court of Justice of the EU has **unlimited jurisdiction to review** ^[te905] Commission fining decisions

Enforcement Timeline ^[tn400]: Phased Application of the AI Act

The AI Act entered into force on 1 August 2024 and has been **in effect from** ^[te103] that date, applying in phases over a three-year period ^[1]. Each phase brings additional provisions into effect, giving affected organizations time to achieve compliance.

- **2 February 2025** [tn401]: Prohibitions on unacceptable-risk AI practices (Chapter II, Art. 5) take effect. AI literacy obligations (Art. 4) require providers and deployers to ensure sufficient AI competence among their staff. This phase is **followed by** [te401] the GPAI rules in August 2025
- **2 May 2025** [tn910]: Deadline for GPAI codes of practice (Art. 56). The AI Office must publish final codes of practice that GPAI model providers can rely on to demonstrate compliance with their obligations. This is **followed by** [te908] the GPAI model obligations in August 2025
- **2 August 2025** [tn402]: Rules for general-purpose AI (GPAI) models apply, including transparency obligations and systemic risk provisions. The governance structure and the penalties framework (including fines up to 35 million EUR or 7% of global turnover) become applicable. This is **followed by** [te402] the high-risk obligations in August 2026
- **2 August 2026** [tn403]: The bulk of the regulation takes effect, with **full obligations from** [te104] this date for high-risk AI systems (Art. 6–27, Art. 40–49), including risk management, data governance, transparency, human oversight requirements, and conformity assessment procedures for providers and deployers. Each Member State must also have an **AI regulatory sandbox** [tn506] (Art. 57) operational by this date. Important: AI systems already placed on the market or put into service before this date are only subject to the Act if they undergo significant changes in their design or intended purpose after 2 August 2026. This is **followed by** [te403] full enforcement in August 2027
- **2 August 2027** [tn404]: **Full enforcement from** [te105] this date: obligations apply to AI systems that are components of large-scale EU IT systems in the areas of freedom, security, and justice listed in Annex X (including SIS, VIS, Eurodac, ETIAS, ECRIS-TCN)

References

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L series. [\[Source\]](#)
2. Corrigendum to Regulation (EU) 2024/1689 of 13 June 2024 (AI Act). *Official Journal of the European Union*, L 2024/1689. [\[Source\]](#)
3. European Commission. AI Act — Shaping Europe's digital future: regulatory framework for AI. [\[Source\]](#)
4. European AI Office. General-Purpose AI Code of Practice — First Draft. [\[Source\]](#)
5. EUR-Lex. Annex III to Regulation (EU) 2024/1689: High-Risk AI Systems referred to in Article 6(2). [\[Source\]](#)

Frequently Asked Questions

Q: What is the EU AI Act?

A: The **EU AI Act** [\[tn500\]](#) (Regulation (EU) 2024/1689) is the world's first comprehensive legal framework specifically regulating **artificial intelligence** [\[tn501\]](#). It establishes harmonised rules for the development, **placing on the market** [\[tn504\]](#), putting into service, and use of AI systems across the European Union, following a risk-based approach that ranges from **prohibited practices** [\[tn100\]](#) to minimal-risk transparency requirements.

Q: Who does the AI Act apply to?

A: The AI Act applies to **providers** [\[tn502\]](#) (who develop or commission AI systems and place them on the market under their own name or trademark), **deployers** [\[tn503\]](#) (who use AI systems under their authority), importers, distributors, and product manufacturers. It covers entities established in the EU, but also those outside the EU if the output of their AI system is used within the Union.

Q: What AI practices are banned under the Act?

A: Article 5 bans eight categories of AI practices deemed to pose unacceptable risk, including **subliminal manipulation** [\[tn101\]](#), **vulnerability exploitation** [\[tn102\]](#), **social scoring** [\[tn103\]](#), **predictive policing** [\[tn104\]](#) based solely on profiling, **untargeted facial scraping** [\[tn105\]](#), **emotion inference in workplace and education** [\[tn106\]](#),

discriminatory **biometric categorization** [tn107], and **real-time remote biometric identification** [tn108] in public spaces for law enforcement (with narrow exceptions).

Q: What are high-risk AI systems and what requirements must they meet?

A: **High-risk** [tn200] AI systems are those used in sensitive domains listed in Annex III, such as **biometrics** [tn201], **critical infrastructure** [tn202], **education** [tn203], and **employment** [tn204]. They must comply with requirements including a **risk management system** [tn301], **data governance** [tn302], **technical documentation** [tn303], **transparency** [tn305], and **human oversight** [tn306], and must undergo **conformity assessment** [tn505] before market placement.

Q: How do I determine if my AI system is high-risk?

A: Follow this decision tree. First, check if your AI system is a safety component of a product covered by Union harmonisation legislation in Annex I (such as medical devices, machinery, or vehicles) — if yes, and the product requires third-party conformity assessment, your system is **high-risk** [tn200]. Second, check if your system falls within one of the eight domains in Annex III. If it does, check whether the **high-risk derogation** [tn209] (Art. 6(3)) applies: the system is NOT high-risk if it performs a narrow procedural task, improves a previously completed human activity, detects decision-making patterns without replacing human judgment, or performs a preparatory assessment task — unless it profiles natural persons, in which case it is always high-risk. If none of these derogations apply, your system is high-risk and must comply with all **compliance requirements** [tn300] in Art. 8–15. Document your classification reasoning and register in the **EU database** [tn908] (Art. 49).

Q: When does the AI Act take effect?

A: The AI Act applies in phases: **prohibited practices** [tn100] and AI literacy from **2 February 2025** [tn401]; GPAI model rules and governance from **2 August 2025** [tn402]; **high-risk** [tn200] AI system obligations from **2 August 2026** [tn403]; and full enforcement for large-scale IT system components from **2 August 2027** [tn404].

Q: What are the penalties for non-compliance?

A: Penalties are tiered by violation severity. The most serious infringements (involving **prohibited practices** [tn100]) can result in fines up to 35 million EUR or 7% of total worldwide annual turnover, whichever is higher. Violations of **high-risk** [tn200] system obligations carry fines up to 15 million EUR or 3% of turnover. Supplying incorrect information to authorities can lead to fines up to 7.5 million EUR or 1% of turnover. Reduced caps apply to SMEs and startups.

Q: What does the AI Act mean for me as an employee?

A: The AI Act affects you in several ways depending on your role. First, **AI literacy** [tn801] (Art. 4) requires your organization to ensure you have sufficient understanding of AI systems you work with — this obligation applies from 2 February 2025. Second, if you are involved in developing AI (**provider obligations** [tn701]), you must ensure compliance with requirements like risk management, data governance, and documentation. If you deploy AI tools (**deployer obligations** [tn702]), you must follow instructions for use, assign human oversight, and inform affected persons. Even using a chatbot or AI content generator triggers **limited risk transparency** [tn550] obligations.

Q: What are general-purpose AI models and why do they matter?

A: **General-purpose AI models** [tn600] (GPAI) are large-scale models trained on broad data that can perform a wide range of tasks — such as large language models and multimodal foundation models. The AI Act regulates these separately because they **underpin** [te601] many downstream AI systems. If your organization uses or integrates a GPAI model, you should understand the provider's obligations regarding technical documentation, copyright compliance, and — if the model has **systemic risk** [tn601] — additional evaluation and cybersecurity requirements.

Q: Who enforces the AI Act and who should I contact with questions?

A: Enforcement operates at two levels. The **EU AI Office** [tn901] (Art. 64) supervises GPAI model providers at the Union level. At the national level, each Member State designates **national competent authorities** [tn903] (Art. 70) — specifically market surveillance authorities — that carry out inspections, handle complaints, and impose penalties. The **European AI Board** [tn902] (Art. 65) coordinates between these national authorities. Within your organization, your compliance, legal, or AI governance team should be your first point of contact.