

# RISK PROFILING IN THE EU: FUNDAMENTAL RIGHTS AND THE MISALIGNMENT OF THE GDPR AND THE AI ACT

Privacy and Non-Discrimination subcommittee

*Marijn De Ruiter*

*Francesca Carantoni*

*Line Kroon*

*Syahna Khalisa Sasyifa*

Young  
— Meijers  
Committee

This comment analyses how EU law regulates profiling, with a specific focus on the recently enacted Regulation (EU) 2024/1689 (the AI Act).

The coexistence of different legislative instruments produces regulatory gaps and inconsistencies. Systems capable of inferring sensitive attributes may escape high-risk classification under the AI Act, while GDPR safeguards often remain unenforceable in practice because of opacity, delegation of responsibility, and weak human oversight. The Law Enforcement Directive adds another asymmetry by permitting broad restrictions on data-subject rights where AI is used for law-enforcement purposes.

Accordingly, this comment asks: How could the AI Act be strengthened to address discriminatory profiling risks more effectively, considering lessons learnt from the GDPR and other relevant EU instruments?

To answer this, the comment first considers the divergence between human-oversight duties and right to meaningful human review. Next, it discusses general-purpose AI, the legal treatment of inferred sensitive data, and the allocation of the burden of proof in discrimination cases.

It moves to consider profiling under the Law Enforcement Directive and concludes by creating a coherent roadmap for aligning EU digital-governance law with the Charter of Fundamental Rights of the European Union rights to equality and effective judicial protection.

It is recommended that system-level oversight under the AI Act should include explicit safeguards ensuring that it complements, and does not dilute, core rights under the GDPR and the Law Enforcement Directive.

In particular, the framework should guarantee meaningful human review, robust transparency, effective access to remedies, and judicial oversight, especially where automated profiling produces legal or similarly significant effects on individuals. Furthermore, regulatory gaps between the AI Act, the GDPR, the Law Enforcement Directive, and EU equality law should be formally closed.

**RISK PROFILING IN THE EU: FUNDAMENTAL RIGHTS AND THE  
MISALIGNMENT OF THE GDPR AND THE AI ACT**

1. Introduction
2. Issues of Technology Governance
  - 2.1 Introduction
  - 2.2 Human Oversight and the Limits of Safeguards
    - 2.2.1 Recommendations
  - 2.3 Enforcement Architecture
    - 2.3.1 Fragmented Institutional Competences
    - 2.3.2 Enforcement Coordination Mechanisms
    - 2.3.3 Recommendations
  - 2.4 Law-Enforcement Context (LED) and Profiling
    - 2.4.1 Legal Framework and Scope
    - 2.4.2 Interaction with the AI Act
    - 2.4.3 Case Law and Supervisory Guidance
    - 2.4.4 Recommendations
  - 2.5 Conclusion on Issues of Technology Governance
3. Technical Limitations
  - 3.1 General Purpose AI
    - 3.1.1. The Structural Link with Other Sections
    - 3.1.2 Profiling Through GPAI: A Regulatory Blind Spot
    - 3.1.3 Systemic-Risk Thresholds and the Equality Gap
    - 3.1.4 The Open-Source Exemption

3.1.5 Connecting GPAI Regulation to the Burden-of-Proof Problem

3.1.6 Recommendations

3.2 Inferred Sensitive Data

3.2.1 Ambiguities under Article 9 GDPR

3.2.2 Procedural Focus Under Article 10 AI Act

3.2.3 Structural Tensions and Ambiguities Between the Two Regimes

3.2.4 Recommendations

3.3 Burden of Proof

3.3.1 Recommendations

3.4 Conclusion on Technical Limitations

4. Concluding Remarks

## 1. Introduction

Artificial intelligence (“AI”) technologies increasingly influence decisions about individuals’ access to credit, employment, welfare, and policing.<sup>1</sup> These decisions frequently rely on profiling: the automated analysis of personal or behavioural data to infer characteristics or predict future behaviour.<sup>2</sup> When such profiles involve or generate information revealing special-category data, such as race, religion, or sexual orientation, the practice constitutes the processing of inferred sensitive data (“ISD”) under Regulation (EU) 2016/679 (“General Data Protection Regulation”).<sup>3</sup>

This comment analyses how EU law regulates profiling, with a specific focus on the recently enacted Regulation (EU) 2024/1689 (“AI Act”). The analysis focuses on improving the fundamental rights standards within this instrument and in light of other legislative instruments, including the GDPR and Directive (EU) 2016/680 (“Law Enforcement Directive”).<sup>4</sup> These instruments pursue overlapping aims - protecting fundamental rights (GDPR) and ensuring trustworthy AI (AI Act) - but operate through different legal mechanisms. The GDPR and LED safeguard individuals’ ex post rights to information, contestation, and redress. The AI Act, by contrast, imposes ex ante obligations on providers and deployers to ensure compliance and safety of AI systems.<sup>5</sup>

Yet the coexistence of these regimes produces regulatory gaps and inconsistencies, particularly in the context of profiling. Systems capable of inferring sensitive attributes may escape high-risk classification under the AI Act, while GDPR safeguards often remain unenforceable in practice because of opacity, delegation of responsibility, and weak human oversight.<sup>6</sup> The LED

---

<sup>1</sup>E. Aizenberg and J. van den Hoven “Desingning for Human Rights in AI” (2020) <<https://arxiv.org/abs/2005.04949>>.

<sup>2</sup> Ibid.

<sup>3</sup> Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (GDPR) [2016] OJ L119/1, arts 4(4) and 9(1).

<sup>4</sup> Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA (LED), 2016 OJ L119/89.

<sup>5</sup> Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) COM (2021) 206 final; see also Council Regulation (EU) 2024/1689 adopting the AI Act OJ L 2024/203.

<sup>6</sup> G. Lazcoz and P. de Hert, ‘Humans in de GDPR and AIA Governance of Automated and Algorithmic Systems. Essential Pre-Requisites Against Abdicating Responsibilities’ (2022) Brussels Privacy Hub Vol. 8, nr. 32.

adds another asymmetry by permitting broad restrictions on data-subject rights where AI is used for law-enforcement purposes.<sup>7</sup>

Accordingly, this comment asks: How could the AI Act be strengthened to address discriminatory profiling risks more effectively, considering lessons learnt from the GDPR and other relevant EU instruments?

The analysis proceeds in three parts: Section 2 examines issues in the governance of technology, namely, the divergence between human-oversight duties and right to meaningful human review, and the fragmented enforcement architecture among competent authorities. Section 3 addresses technological and substantive issues, focusing on general-purpose AI (“GPAI”), the legal treatment of inferred sensitive data, and the allocation of the burden of proof in discrimination cases. Section 4 briefly considers profiling under the LED, highlighting information-rights restrictions and proposing minimum-transparency safeguards. A concluding section sets out our recommendations, creating a coherent roadmap for aligning EU digital-governance law with the Charter of Fundamental Rights of the European Union (“the Charter”) rights to equality and effective judicial protection.<sup>8</sup>

---

<sup>7</sup> Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties (LED) [2016] OJ L119/89, art 13(3).

<sup>8</sup> Charter of Fundamental Rights of the European Union [2012] OJ C326/391, arts 21 and 47. See also EDPB and EDPS, ‘Joint Opinion 5/2021 on the proposal for a Regulation laying down harmonised rules on Artificial Intelligence’ (18 June 2021) para 13.

## 2. Issues of Technology Governance

### 2.1 Introduction

This section analyses the ambiguities in governance structures between the several EU instruments that regulate obligations related to profiling. The analysis is tailored around the principle of human oversight (section 2.2) and the enforcement architecture (2.3). We conclude this section by drawing your attention to our recommendations.

### 2.2 Human Oversight and the Limits of Safeguards

In its 2021 proposal for the AI Act, the European Commission opined that:

*AI should be a tool for people and be a force for good in society with the ultimate aim of increasing human well-being. Rules for AI available in the Union market or otherwise affecting people in the Union should therefore be human-centric, so that people can trust that the technology is used in a way that is safe and compliant with the law, including the respect of fundamental rights.<sup>9</sup>*

The Commission acknowledged the merits of human oversight and as a result, the AI Act includes requirements for a human-centric approach. On 2 August 2026 the full AI Act will be applicable, including its Article 14. In line with this article, providers and deployers of high-risk AI systems must comply with the human oversight principle, as concerns high-risk AI systems.

### Divergent Objectives and Scopes

While both the AI Act and the GDPR aim to ensure human agency in the digital domain, they do so in different ways:

- Article 22 GDPR protects individuals via a rights-based approach, securing access to human intervention and contestation
- Article 14 AI Act by contrast, establishes the technical and organisational human oversight obligation at system-level for high-risk AI systems

Unlike Article 22 GDPR, which grants individuals the right to not be subject to fully automated decision-making and profiling decisions, Article 14 imposes an obligation on providers to ensure humans design and develop their use. In other words, while the GDPR safeguards the

---

<sup>9</sup> Commission, ‘Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts’ SEC (2021) 167 final, SWD (2021) 84 final, SWD (2021) 85 final, Explanatory Memorandum, p. 1.

individual by ensuring the right to human intervention and contestation, the AI regulates the system by requiring that human oversight is structurally embedded.

The rationale between both provisions is similar: to create a human-centric digital world. In short, this entails human control over digital tools, where the respect for fundamental rights and principles is entrenched.<sup>10</sup> However, in the practical application of these legal instruments, there are certain discrepancies between them that create an ambiguous climate for compliance.<sup>11</sup>

Firstly, the principles of ‘human oversight’ in the AI Act and ‘human intervention’ in the GDPR are not interchangeable.<sup>12</sup> While Article 22 GDPR ensures that the data subject can get access to meaningful information about the logic of the decision or the profiling and allows it to contest its outcome. The human intervention measures in Article 22 GDPR come with a stricter regime than the measures provided in Article 14 AI Act. Consequently, being compliant with Article 14 AI Act does not ensure compliance with Article 22 GDPR due to divergences in scope and definition.

Additionally, the obligation contained in Article 14 AI Act is solely applicable in relation to high-risk AI systems, whereas Article 22 GDPR is applicable to any automated decision, or automated profiling with significant effects for the data subject, unless a paragraph 2 exemption applies. Therefore, the application scope of Article 22 GDPR is significantly broader, which creates compliance challenges in relation to usage of medium risk and low risk AI systems that perform risk profiling or automated decision tasks, since they may assume that no measures are required due to the fact that these systems fall outside the scope of Article 14 AI Act.

Lastly, the compliance responsibilities between Article 14 AI Act and Article 22 GDPR are distributed differently. Article 14 AI Act places responsibilities on providers and deployers of the AI systems, whereas the GDPR focuses merely on data controllers and processors. These actors might or might not be the same entity. Moreover, the legal duties related to these responsibilities do not resemble each other, which complicates accountability in profiling contexts involving multiple actors.

Thus, while both Article 14 AI Act and Article 22 GDPR aim to safeguard human agency in automated decision-making, they do so through distinct legal mechanisms. Article 22 GDPR establishes the right for the individual not to be subject to automated decisions, whereas Article

---

<sup>10</sup> Ibid.

<sup>11</sup> L. Colonna, ‘Exploring the Relationship Between Article 22 of the General Data Protection Regulation and Article 14 of the Proposed AI Act’ (2022) Stockholm Faculty of Law University Research Paper no. 124.

<sup>12</sup> Ibid.

14 AI Act imposes an obligation on providers to design-high risk AI systems in a way that ensures they remain subject to human oversight.

### 2.2.1 Recommendations

1. We suggest inserting an explicit non-prejudice clause in Article 14 AI Act cross-referring to Articles 15(1)(h) and 22 GDPR, to ensure that system-level oversight complements, rather than displaces, individual rights.
2. We recommend defining meaningful human review with reference to guidance by the AI Office and European Data Protection Board. Within this definition, it ought to be stated that human review requires reviewers to possess decision-making authority, contextual expertise, and the power to override automated outputs.
3. We recommend that joint interpretative guidelines be created by the European Commission for authorities applying the AI Act, GDPR, and LED, clarifying oversight obligations in cross-regime contexts (e.g., public–private data sharing).

### 2.3 Enforcement Architecture

Even where human-oversight mechanisms are sound, the enforcement architecture of EU digital-governance law remains fragmented. Supervisory competence is divided between data-protection authorities (“DPAs”) enforcing the GDPR and LED, market-surveillance authorities (“MSAs”) designated under the AI Act, and soon the European AI Office responsible for coordination, systemic-risk supervision, and guidance.<sup>13</sup> This institutional pluralism risks diverging monitoring, as each authority monitors only a segment of the AI life-cycle.

#### 2.3.1 Fragmented Institutional Competences

Under Article 59 AI Act, Member States must designate MSAs to monitor compliance by providers and deployers of high-risk systems, while the AI Office ensures consistency at EU level. These bodies, however, lack competence to adjudicate individual rights claims, which fall within the remit of DPAs and national courts under the GDPR.<sup>14</sup> Conversely, DPAs have no direct powers to evaluate conformity-assessment files or risk-management documentation prepared under Articles 9–10 AI Act. The result is an asymmetric enforcement model: systemic-risk supervision detached from rights enforcement.

This separation becomes critical in profiling cases. A deployer may rely on an AI-based scoring tool that infers sensitive attributes from non-sensitive data. The provider’s obligations, risk

---

<sup>13</sup> AI Act arts 59–64; Regulation (EU) 2016/679 (GDPR) arts 51–58; Directive (EU) 2016/680 (LED) arts 41–44.

<sup>14</sup> AI Act arts 9–10; GDPR art 57(1)(a)–(f).

management, transparency and data-quality documentation, fall under the AI Act and MSA control, whereas the deployer’s obligations to provide access, rectification, and human review lie under the GDPR and DPA oversight. No single authority follows the case end-to-end, meaning that rights violations can slip through the cracks.<sup>15</sup>

### 2.3.2 Enforcement Coordination Mechanisms

Despite the existence of several coordination routes, these routes remain non-binding in nature. Recital 110 AI Act encourages cooperation between MSAs and DPAs where AI systems involve personal-data processing, yet the Regulation stops short of imposing mandatory procedures or timelines. The European Data Protection Board (“EDPB”) has repeatedly urged for the creation of joint investigation mechanisms, noting that “fragmented supervision undermines the effective enforcement of fundamental-rights safeguards.”<sup>16</sup> Without formal channels for evidence-sharing and mutual recognition of findings, authorities duplicate efforts or defer responsibility.

### 2.3.3 Recommendations

To operationalise oversight across the full AI lifecycle, we recommend a three-layered model approach:

1. Vertical coordination: the AI Office coordinates national MSAs that subsequently coordinate with DPAs. We recommend that the AI Office issues binding guidelines on cooperation and takes an oversight role in relation to MSAs to transmit inspection outcomes to DPAs when personal-data processing is identified. Additionally, we recommend the establishment of a formal Memorandum of Understanding between the AI Office, MSAs, DPAs, and equality bodies, to ensure mutual reliance on each other’s findings in profiling and discrimination cases.
2. We recommend greater horizontal coordination between DPAs and national equity bodies. Such collaboration would address non-discrimination harms, complementing data-protection enforcement through shared evidence standards and disclosure orders.
3. We recommend that a one-stop complaint interface be created: A joint online portal where individuals can submit profiling-related grievances, automatically redirecting them to the competent authority while preserving cross-reference for follow-up.<sup>17</sup>

---

<sup>15</sup> EDPB, *Guidelines 3/2019 on processing of personal data through video devices* (adopted 29 January 2020) paras 8–10 (noting limits of sectoral supervision).

<sup>16</sup> EDPB and EDPS, *Joint Opinion 5/2021 on the proposal for a Regulation laying down harmonised rules on Artificial Intelligence* (18 June 2021) paras 16–18.

<sup>17</sup> Inspired by the *European Consumer Protection Cooperation (CPC)* network model under Regulation (EU) 2017/2394 OJ L 345/1, which mandates information-sharing and mutual assistance among national authorities.

4. We recommend the creation of an EU-level task force on profiling and inferred sensitive data, hosted by the AI Office with DPA participation, to develop shared inspection templates and risk indicators.
5. We recommend the creation of a provision within the AI Act to require risk-management documentation under Article 10 AI Act be made accessible by providers, upon request, to DPAs and equality bodies for cross-assessment of discrimination claims.

## 2.4 Law-Enforcement Context (LED) and Profiling

Profiling and AI-assisted risk scoring are not confined to the private sector. They are increasingly deployed by law-enforcement authorities (“LEAs”) for predictive policing, identity verification, and criminal-risk assessment.<sup>18</sup> In this sphere, the applicable framework consists of the Law Enforcement Directive (“LED”), which largely mirrors the GDPR but incorporates broad restrictions on transparency and access, and the AI Act. These limitations can significantly weaken safeguards against discriminatory profiling.

### 2.4.1 Legal Framework and Scope

The LED governs the processing of personal data by competent authorities for law-enforcement purposes. Articles 12–14 set out the right to information and access, but Article 13(3) allows Member States to restrict these rights whenever their exercise would prejudice investigations, public security, or national security. Because these exceptions are drafted in general terms and do not require proportionality testing or time limits, they can easily become the default rule in sensitive policing contexts.<sup>19</sup>

AI-based profiling used in law-enforcement operations often combines data from multiple sources, including privately developed general-purpose models. Once a LEA procures or re-uses such systems, oversight responsibility becomes blurred. The AI Act excludes most purely public-security activities from its scope (Article 2(3)),<sup>20</sup> while DPAs’ competence under the LED is limited to ensuring compliance by competent authorities within their own Member State. The result is a patchwork of national oversight practices and an absence of EU-level monitoring regarding how AI profiling affects fundamental rights in policing.

---

<sup>18</sup> R.A. Berk, ‘Artificial Intelligence, Predictive Policing, and Risk Assessment for Law Enforcement’ (2021) *Annual Review Criminology* vol. 2021/4.

<sup>19</sup> Directive (EU) 2016/680 (LED) arts 12–14, 13(3). See also EDPB, *Guidelines 06/2022 on the practical implementation of the LED* (Adopted 28 June 2022) paras 12–15.

<sup>20</sup> AI Act art 2(3); recital 12.

## 2.4.2 Interaction with the AI Act

Although the AI Act does not apply directly to national-security or defence uses, it does cover providers placing AI systems on the market, including those later acquired by LEAs. Providers must therefore comply with the AI Act’s risk-management and transparency obligations, even if deployers operate under the LED. However, Articles 10 and 14 AI Act impose no duties on LEAs themselves, meaning that once the system is in use, the individual’s ability to obtain information or contest decisions depends entirely on the LED’s weaker rights regime. This creates a “regulatory cliff-edge”: strong ex-ante duties for private suppliers but minimal ex-post guarantees for individuals affected by public-sector profiling.<sup>21</sup>

## 2.4.3 Case Law and Supervisory Guidance

While the Court of Justice of the European Union has not yet interpreted Article 13(3) LED directly, its recent *SCHUFA* and *CK v Dun & Bradstreet* judgments demonstrate a general commitment to ensuring transparency in automated decision-making.<sup>22</sup> These rulings can guide analogous interpretation of the LED: restrictions on access should never nullify the essence of the right to meaningful information under Article 8(2) Charter of Fundamental Rights of the EU (“Charter”). The European Data Protection Board (EDPB) Guidelines 05/2022 on facial-recognition technology likewise stress that any limitation on data-subject rights must be strictly necessary and subject to independent oversight.<sup>23</sup>

## 2.4.4 Recommendations

1. We recommend the incorporation of minimum-information thresholds. By this, we encourage the harmonisation of national implementation with respect to Article 13(3) LED to guarantee that individuals receive at least general notice of automated processing and the categories of data used, even where detailed disclosure is restricted.
2. We further recommend that, once an investigation risk has ceased, affected persons are mandatorily informed that profiling took place and of their right to seek judicial review.
3. Additionally, we encourage measures to increase cross-regime cooperation. We encourage cooperation between DPAs supervising LED processing and MSAs enforcing the AI Act when the same system is used for both public- and private-sector purposes.

---

<sup>21</sup>ibid arts 10 and 14; EDPB–EDPS *Joint Opinion 5/2021 on the proposal for a Regulation laying down harmonised rules on Artificial Intelligence* (18 June 2021) para 19.

<sup>22</sup> *SCHUFA Holding AG* Case C-634/21 EU:C:2023:1032 paras 64–69; *CK v Dun & Bradstreet Austria GmbH* Case C-634/21 EU:C:2025:146 paras 51–55.

<sup>23</sup> EDPB, *Guidelines 05/2022 on the use of facial-recognition technology by law-enforcement authorities* (28 June 2022) paras 17–23.

4. Furthermore, at the EU Level, greater oversight could benefit the way in which these regulations are applied. For instance, empowering the AI Office, together with the European Union Agency for Fundamental Rights (“FRA”), to publish periodic reports on the use of AI-based profiling by LEAs and its impact on equality and non-discrimination could improve the situation.

## 2.5 Conclusion on Issues of Technology Governance

The examination of human oversight and enforcement architecture reveals that the current configuration of EU digital governance, though ambitious in scope, remains structurally fragmented and normatively inconsistent. Both the AI Act and the GDPR aspire to ensure a human-centric digital ecosystem that respects fundamental rights, yet they do so through divergent mechanisms: the GDPR safeguards individuals through rights of human intervention and contestation, while the AI Act imposes system-level obligations of human oversight on providers and deployers. This duality generates interpretive and practical ambiguities, particularly concerning who is responsible for ensuring human oversight and how compliance can be consistently verified across overlapping regulatory regimes.

Such ambiguity undermines the effectiveness of safeguards against discriminatory profiling. While the GDPR’s Article 22 confers a direct, enforceable right on individuals, Article 14 of the AI Act offers only procedural guarantees that may not translate into meaningful redress. The differing scopes of application: individual rights under the GDPR versus systemic obligations under the AI Act – further fragment oversight, especially where medium- and low-risk AI systems still perform profiling functions with significant social impact.

The problem is compounded by an enforcement architecture divided between Data Protection Authorities (“DPAs”), Market Surveillance Authorities (“MSAs”), and the forthcoming European AI Office. Each operates within its own institutional and legal silo: DPAs handle rights-based claims, MSAs oversee technical conformity, and the AI Office provides strategic coordination. However, no authority has an end-to-end view of profiling practices across the AI lifecycle. This institutional pluralism creates enforcement blind spots where discriminatory outcomes may persist undetected or unremedied.

To effectively address profiling-related discrimination, we suggest structuring the AI Act as a bridge between the normative gap (between system-level oversight and individual rights) and the institutional gap (between fragmented enforcement bodies). For strengthening the AI Act, we recommend three essential measures:

1. Normative alignment: explicitly link Article 14 AI Act with Article 22 GDPR by clarifying that compliance with human oversight obligations does not prejudice, and indeed must complement, individuals' rights to human intervention and contestation.
2. Integrated enforcement: establish binding coordination mechanisms, such as a formal MoU or legislation, between the AI Office, DPAs, MSAs, and equality bodies to ensure coherent investigation and enforcement of profiling cases.
3. Transparency and accountability: mandate reciprocal access to risk-management documentation and profiling-related audit reports for all relevant authorities, enabling cross-regime supervision and informed redress for individuals.

Such alignment ensures that the AI Act not only mitigates systemic risks but also strengthens protection against discriminatory profiling in a manner consistent with the fundamental-rights ethos of the GDPR and related instruments.

### 3. Technical Limitations

The technical design of AI systems determines the practical effectiveness of legal safeguards. This section addresses three substantive aspects that directly affect profiling and the generation of inferred sensitive data (“ISD”): (1) general-purpose AI (“GPAI”) systems, (2) the legal treatment of inferred sensitive data under EU law, and (3) the allocation of the burden of proof in discrimination disputes.

#### 3.1 General Purpose AI

##### 3.1.1. The Structural Link with Other Sections

General-purpose AI (GPAI) refers to AI systems that are designed to handle a wide range of tasks and can be adapted for different uses, such as recruitment, credit scoring, or law enforcement. These models are flexible and can be fine-tuned to perform various functions depending on the specific application.<sup>24</sup>

The risks associated with GPAI-models exacerbate. While inferred sensitive data and human oversight focus on specific instances of profiling and control, GPAI serves as the broader, systemic foundation where these risks often originate and proliferate. The opaque nature of GPAI models, combined with the lack of access to their internal data and adjustment parameters, means that individuals often cannot access the necessary information to establish a prima facie case of discrimination, thereby exacerbating the burden of proof (Section 3).<sup>25</sup>

##### 3.1.2 Profiling Through GPAI: A Regulatory Blind Spot

GPAI models are large, versatile systems that can be fine-tuned for diverse applications in employment, credit, welfare or law enforcement. These are regulated under Chapter V AI Act.<sup>26</sup> However, the data protection and non-discrimination safeguards in the GDPR and the Law Enforcement Directive (LED) only apply when such models are deployed in context-specific, “high-risk” uses.<sup>27</sup> This creates a temporal and normative gap: algorithmic profiling may already occur at the model-training or fine-tuning stage, before any “use case” is classified

---

<sup>24</sup> EU Regulation 2024/1689, Article 3(63).

<sup>25</sup> Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 International Data Privacy Law 76 (2017).

<sup>26</sup> European Parliament and Council, *Regulation (EU) 2024/... laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, Chapter V (General-Purpose AI models), Official Journal of the European Union (2024).

<sup>27</sup> Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 International Data Privacy Law 76 (2017).

as high risk. In practice, this undermines Article 9 GDPR's prohibition on processing special categories of data and the Article 22 GDPR right not to be subject to solely automated decisions, since both are triggered too late in the AI lifecycle.<sup>28</sup>

Under the LED, public authorities using GPAI for law enforcement must ensure fairness and non-discrimination (Arts 4–8 LED). However, private developers are not subject to the same requirements, creating an asymmetry. This allows risks to fundamental rights, such as discrimination, to go unchecked when AI is used by private actors.<sup>29</sup>

### 3.1.3 Systemic-Risk Thresholds and the Equality Gap

Article 52 and Annex XIII of the AI Act define systemic-risk GPAI primarily based on their technical capacity, such as requiring more than approximately  $10^{25}$  floating-point operations (FLOPs), a measure of computational power. These models, due to their large-scale processing capabilities, are subject to stricter safeguards. However, smaller models, which don't meet the computational scale threshold, often fall outside this perimeter, leaving potential equality risks unregulated

☞ While this criterion aims to capture models with transformative potential (such as frontier generative systems), it misses equality-related risk. Smaller or open-source models can produce equally serious discriminatory inferences, e.g. by inferring gender or ethnicity from language patterns in job applications or welfare assessments. These harms are independent of computational size and thus remain outside the “systemic-risk” perimeter. The focus on FLOPs, though convenient for technical classification, fails to address fundamental-rights impacts, contrary to Articles 21 and 47 CFREU.<sup>30</sup>

### 3.1.4 The Open-Source Exemption

Article 52c(2) AI Act exempts open-source GPAI models from many obligations unless they meet systemic-risk criteria. While open distribution fosters innovation and transparency, it creates a regulatory vacuum. Open-source models can be fine-tuned to perform profiling in sensitive contexts (recruitment, policing, migration) without any requirement to assess equality impacts or data-protection compliance. This exemption thus undermines both the AI Act's

---

<sup>28</sup> Article 29 Data Protection Working Party, *Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679*, WP251 rev.01 (2018).

<sup>29</sup> Directive (EU) 2016/680 of the European Parliament and of the Council (Law Enforcement Directive), Arts 4–8.

<sup>30</sup> Frederik Zuiderveen Borgesius, *Discrimination, Artificial Intelligence, and Algorithmic Decision-Making*, Council of Europe Study (2018).

protective purpose and the GDPR’s guarantees of lawful, fair, and transparent processing (Art. 5 GDPR).<sup>31</sup>

### 3.1.5 Connecting GPAI Regulation to the Burden-of-Proof Problem

Because GPAI models are often opaque and re-used by multiple downstream actors, victims of discrimination face an impossible evidentiary burden. They lack access to training data, fine-tuning parameters, or internal documentation necessary to establish a prima facie case under the Racial Equality Directive or the Gender Equality Directive. This extends the “black-box” challenge identified earlier: without procedural transparency obligations at the model-level, the burden of proof cannot realistically shift.<sup>32</sup>

### 3.1.6 Recommendations

1. We recommend the introduction of an Equality-Risk Designation Pathway: The Commission and AI Office should be empowered to designate GPAI as “systemic” not only by scale (FLOPs) but also where there is evidence of repeated or potential equality-related harm (e.g. proxy variables, biased outputs).
2. We recommend the integration of the GDPR and AI Act Compliance: Any GPAI model capable of inferring special-category data under Article 9 GDPR should automatically trigger high-risk obligations under Article 10 AI Act, ensuring consistency between fundamental-rights and market-risk logics.
3. Mandatory Coordination Mechanism: Establish formal cooperation between Data Protection Authorities (DPAs), the AI Office, and national equality bodies to exchange model-level documentation and enforce Articles 21 and 47 CFR effectively.
4. Clarify Technical Terminology: Replace opaque measures like “10 FLOPs” with explanatory guidance (e.g., “equivalent to training on hundreds of high-end GPUs for several weeks”) to ensure accessibility to policymakers and legal practitioners. This recommendation aims to address the discriminatory risks inherent in the profiling processes of GPAI by ensuring that systems causing inequality are adequately regulated.

## 3.2 Inferred Sensitive Data

Inferred sensitive data are attributes derived or predicted from other data that reveal protected characteristics, such as health status, ethnicity, sexuality, or political beliefs. These traits are

---

<sup>31</sup> European Commission, *Explanatory Memorandum to the Proposal for a Regulation laying down harmonised rules on Artificial Intelligence*, COM(2021) 206 final; see also Open Future, *Open-Source AI and the EU AI Act* (policy brief, 2023).

<sup>32</sup> Solon Barocas & Andrew D. Selbst, *Big Data’s Disparate Impact*, 104 California Law Review 671 (2016).

not explicitly provided by individuals but deduced through algorithmic analysis of seemingly unrelated information.<sup>33</sup> For instance, ethnicity may be inferred from name patterns or geographic location, and sexual orientation from browsing history or social-media behaviour.<sup>34</sup>

The legal treatment of inferred sensitive data in EU law sits at the intersection of Article 9 GDPR and Article 10 AI Act, yet both provisions leave significant regulatory gaps.

### 3.2.1 Ambiguities under Article 9 GDPR

Although Article 9(1) GDPR prohibits the processing of personal data revealing special categories such as racial or ethnic origin, political opinions, or health status, its wording primarily targets explicit data collection, leaving ambiguity about algorithmic inferences that reveal such attributes indirectly. Controllers frequently exploit this uncertainty, claiming that inferences are “mere guesses” or statistical abstractions rather than “processing” of personal data.<sup>35</sup>

However, this position has been rejected by the European Data Protection Board's (EDPB) Guidelines 05/2020 on Consent, which clarify that inferences revealing sensitive aspects about an identifiable individual constitute sensitive personal data, even if generated probabilistically or without the data subject's direct disclosure.<sup>36</sup> Despite this clarification, enforcement practice remains inconsistent. Controllers often argue that transient inferences—those used momentarily for targeting or scoring and not stored—fall outside the definition of “processing”.<sup>37</sup> Others claim that pseudonymisation removes inferred data from Article 9's scope. Furthermore, consent for such processing is often obtained through bundled “personalisation” clauses rather than explicit consent as required under Article 9(2)(a), undermining the provision's substantive protection.<sup>38</sup>

---

<sup>33</sup> Wachter S and Mittelstadt B [2018] A right to reasonable inferences: Re-thinking data protection law in the age of big data and ai.

<sup>34</sup> *ibid.*

<sup>35</sup> Gioia G and Lener SM, ‘The Protection of Individuals against Privacy-Invasive and Discriminatory Inferences under European Law: From the General Data Protection Regulation and the Digital Content and Services Directive to the Artificial Intelligence Act’ (2024) 74 *Zbornik Pravnog fakulteta u Zagrebu* 861.

<sup>36</sup> European Data Protection Board, *Guidelines 05/2020 on Consent under Regulation 2016/679* (adopted 4 May 2020) paras 91–93.

<sup>37</sup> Damian George, Kento Reutimann and Aurelia Tamò-Larrioux, ‘GDPR Bypass by Design? Transient Processing of Data under the GDPR’ (2019) 9 *International Data Privacy Law* 285; Wachter (n 18).

<sup>38</sup> Giuseppe Ziccardi, *Consent as Friction: A Critique of GDPR Consent in Digital Markets* (2019) 60 *Boston College Law Review* 1381, 1402–03.

### 3.2.2 Procedural Focus Under Article 10 AI Act

Article 10 AI Act addresses sensitive data through a risk-governance lens rather than a prohibition-based model. Providers of high-risk AI systems must ensure datasets are relevant, representative, free of errors, and sufficiently complete, implementing safeguards to prevent discriminatory outcomes.<sup>39</sup> However, these are procedural obligations, not substantive prohibitions.

In other words, while the AI Act's Article 10 imposes procedural data-quality obligations, it does not replicate the GDPR's substantive prohibition of sensitive-data processing. As long as a provider demonstrates documentation, validation, and mitigation measures, the generation of sensitive inferences remains permissible.<sup>40</sup> Loopholes persist where actors classify their tools as non-AI systems, general-purpose AI, or claim exemptions for law enforcement or migration contexts.<sup>41</sup> Moreover, the prohibition on biometric categorisation in Article 5 AI Act is narrowly confined to certain real-time practices, leaving most offline or indirect inferences—and the profiling they enable—unregulated.<sup>42</sup>

### 3.2.3 Structural Tensions and Ambiguities Between the Two Regimes

Taken together, Article 9 GDPR imposes categorical limits on processing that reveals sensitive attributes but is vulnerable to restrictive interpretation, whereas Article 10 AI Act accepts the generation of such inferences as inevitable, focusing only on procedural risk control.<sup>43</sup> This divergence reflects the instruments' distinct purposes: the GDPR protects individual rights and freedoms in data processing, while the AI Act primarily serves as a market regulation tool designed to manage systemic risks of high-risk AI systems.<sup>44</sup>

In practice, AI developers and deployers must navigate both regimes simultaneously ensuring that algorithmic inferences do not unlawfully reveal special-category data under Article 9 GDPR, while implementing dataset governance and risk-management measures under Article 10 AI Act.<sup>45</sup> Yet the absence of alignment between these frameworks creates a fragmented

---

<sup>39</sup> Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 1689/1, art 10.

<sup>40</sup> Sandra Wachter, Brent Mittelstadt and Chris Russell, 'Why Fairness Cannot Be Automated: Bridging the Gap Between EU Non-Discrimination Law and AI' (2021) 41.

<sup>41</sup> *ibid.*

<sup>42</sup> *ibid.*

<sup>43</sup> Regulation (EU) 2016/679 (General Data Protection Regulation) [2016] OJ L 119/1, art 9; Regulation (EU) 2024/1689 (Artificial Intelligence Act) [2024] OJ L 1689/1, art 10

<sup>44</sup> Sandra Wachter and Brent Mittelstadt, 'A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI' (2019) 2 *Columbia Business Law Review* 494.

<sup>45</sup> Regulation (EU) 2016/679 (GDPR) art 9; Regulation (EU) 2024/1689 (AI Act) art 10;

regulatory environment in which sensitive inferences are neither clearly prohibited nor consistently constrained, allowing profiling practices that exploit these legal ambiguities to persist.<sup>46</sup>

### 3.2.4 Recommendations

1. We recommend clarifying the Scope of Article 9 GDPR: explicitly recognise that algorithmic inferences revealing special-category data constitute “processing”, irrespective of whether such data are stored, pseudonymised, or expressed probabilistically.
2. Additionally, we suggest strengthening consent standards. Clear guidance needs to be provided on what qualifies as explicit consent for inferred sensitive data, ensuring that bundled or generic “personalisation” consents cannot legitimise processing revealing protected characteristics.
3. We propose to reinforce Article 10 AI Act obligations by extending the safeguards under Article 10 to all contexts where sensitive traits are inferred. This includes requiring providers not only to document mitigation measures but also to justify the necessity and proportionality of generating such inferences.
4. We also suggest further integration of the GDPR and AI Act framework. Specifically establishing that any AI system producing or relying on special-category data within the meaning of Article 9 GDPR automatically triggers the high-risk obligations under Article 10 AI Act. This linkage would close existing regulatory gaps and prevent compliance arbitrage across the two regimes.

### 3.3 Burden of Proof

The starting point for the burden of proof is Article 6(2) of the European Convention on Human Rights and Fundamental Freedoms (“ECHR”), stating that “everyone charged with a criminal offence shall be presumed innocent until proved guilty according to law”. This is the essence of the burden of proof in criminal law. In the context of non-discrimination, both Article 8(1) Directive 2000/43/EC (Racial Equality Directive) and Article 19(1) Directive 2006/54/EC (Gender Equality Directive) determine that the burden of proof shifts to the defendant if a prima facie discrimination case is established.<sup>47</sup> The idea behind shifting the burden of proof is to

---

European Data Protection Supervisor (EDPS), *Opinion 5/2021 on the Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)* [2021] OJ C 92/1, paras 17–26

<sup>46</sup> Sandra Wachter, Brent Mittelstadt and Chris Russell, ‘Why Fairness Cannot Be Automated: Bridging the Gap Between EU Non-Discrimination Law and AI’ (2021) 41 *Computer Law & Security Review*; Damian George, Kento Reutimann and Aurelia Tamò-Larrieux, ‘GDPR Bypass by Design? Transient Processing of Data under the GDPR’ (2019) 9 *International Data Privacy Law* 285.

<sup>47</sup> Council Directive 2000/43/EC of 29 June 2000 implementing the principle of equal treatment between persons irrespective of racial or ethnic origin 22 OJ L180 (*Racial Equality Directive*); Directive 2006/54/EC of the

protect victims of discrimination when the case is dependent on factors within the knowledge and control of the defendant.<sup>48</sup> The defendant must then provide evidence that there has been no case of discrimination.

Thus, the claimant should provide sufficient evidence to prove discrimination on the side of the defendant before the burden of proof shifts to the defendant. For example, in the field of employment, for the burden of proof to shift, the worker has to provide a prima facie case that discrimination has occurred (e.g., the women in the company get paid less than the men that do the same job or work of equal value).<sup>49</sup> Once the prima facie case of discrimination has been established, the burden shifts to the employer to provide evidence that disproves that discrimination occurred. Effective access to facts and evidence is crucial to establish a prima facie case of discrimination.

Providing a prima facie threshold to prove discrimination in AI systems is more complicated, as these often operate as black boxes, not providing clear explanations of the decision-making process.<sup>50</sup> The particular characteristics of deep and self-learning AI systems makes it especially difficult for individuals to establish a prima facie case of discrimination.<sup>49</sup> There are various characteristics of AI systems that create specific challenges for the burden of proof to shift, such as lack of causality, opacity, unpredictability, and self and continuous learning capability.<sup>51</sup> As Guo and others point out, most data-driven AI is able to identify correlations rather than causal relations. Due to the lack of independent variables, AI cannot make inferences about causal relations, meaning that the variables that these systems work with do not meet the criteria of independence, and therefore, the system is only able to make inferences based on correlations between the variables.<sup>52</sup> Another challenge is opacity, i.e. the logic behind the decision-making process of AI systems is often too complex to understand, even for experts.<sup>53</sup> This is often also referred to as the black-box effect.<sup>54</sup>

---

European Parliament and of the Council of 5 July 2006 on the implementation of the principle of equal opportunities and equal treatment of men and women in matters of employment and occupation (recast) 23 OJ L204 (*Gender Equality Directive*).

<sup>48</sup> Farkas L and O'Farrell O (European Commission 2015) rep

<sup>49</sup> *ibid.*

<sup>50</sup> Bathaee Y, 'The Artificial Intelligence Black Box and the Failure of Intent and Causation' (2018) 31 *Harvard Journal of Law & Technology* 890; Pedreschi D and others, 'Meaningful Explanations of Black Box AI Decision Systems' (2019) 33 *Proceedings of the AAAI Conference on Artificial Intelligence* 9780.

<sup>51</sup> Fernández Llorca D and others, 'Liability Regimes in the Age of AI: A Use-Case Driven Analysis of the Burden of Proof' (2023) 76 *Journal of Artificial Intelligence Research* 613.

<sup>52</sup> Guo R and others, 'A Survey of Learning Causality with Data' (2020) 53 *ACM Computing Surveys* 1

<sup>53</sup> *Ibid.* 3

<sup>54</sup> Castelvechhi D, 'Can We Open the Black Box of Ai?' (2016) 538 *Nature* 20

Consequently, the application of AI in establishing a risk profile, may jeopardise the objectives of the Racial Equality Directive and Gender Equality Directive respectively. Due to the black box nature of AI systems, subjects could be deprived of the necessary information to establish evidence of a prima facie case of discrimination. As a result, the protection granted by shifting the burden of proof in non-discrimination law could be weakened as it may be challenging to establish a prima facie discrimination case when certain information is not available because of the use of AI. Article 22 GDPR allows the subject to request this information to contest a decision, but this might be challenging if the information simply does not exist due to the AI tool's nature. Therefore, applying AI tools in the context of risk profiling can easily result in a discriminatory outcome which may be challenging to contest.

Article 5(d) AI Act prohibits the use of AI systems specifically for risk assessments in the criminal context. However, this provision does not apply for AI systems that support human risk assessment, nor does it apply in civil context (e.g. non-discrimination law). Considering the above, the AI act does not ensure that the level of protection granted in the Racial Equality Directive and Gender Equality Directive to prevent discrimination is maintained. It neglects to introduce the necessary safeguards to prevent discrimination when AI tools are used for risk profiling in a non-criminal context.

The challenge to shift the burden of proof in cases of discrimination through AI lies in the complexity, opacity and unpredictability of current AI systems. In its judgment, the Court of Justice of the European Union, clarified that individuals are entitled to a meaningful explanation of the principles and logic underlying automated decision-making, such as credit scoring.<sup>55</sup> Yet, even with such safeguards, significant challenges can remain, especially for deep and self-learning AI systems. The complexity that these systems have due to their adaptive and multi-layered nature means that explanations are often too difficult to understand. As a result, there remains an obstacle to access evidence in cases of discrimination via such AI systems.

### 3.3.1 Recommendations

Based on the above, there are various aspects to consider that could potentially minimise the risk of discriminatory outcomes when using AI:

1. We recommend adopting a new provision, Article 5(1) (da), which explicitly prohibits the use of AI systems that generate discriminatory risk profiles of natural persons in non-criminal settings. This provision should specifically cover profiling practices that

---

<sup>55</sup> Case C-203/22, *Ck v Dun & Brandstreet Austria GmbH*

rely on, or derive inferences from, data relates to race, ethnicity and gender, given the heightened risk of structural bias based on these characteristics. As such, it should reduce the discriminatory application of AI systems used for risk profiling and strengthen the legal position of individuals affected.

2. We advise the provision explicitly targets AI systems used for risk assessment, eligibility determination or predictive profiling in domains such as employment, housing, social services and access to public or private goods. Clarifying the scope could enhance legal certainty for AI providers and deployers, while ensuring effective protection for individuals affected by discriminatory algorithmic decision-making.
3. We further recommend that this prohibition of such AI systems be explicitly aligned with Article 8(1) of Directive 2000/43 (racial equality) and Article 19(1) of Directive 2006/54/EC (gender equality). This can promote coherence across EU legal frameworks and ensure reinforcement of the AI Act.

### 3.4 Conclusion on Technical Limitations

The analysis of technical limitations reveals that the AI Act's capacity to prevent discriminatory profiling is weakened not only by regulatory fragmentation but also by how technological design interacts with legal safeguards. Across general-purpose AI, inferred sensitive data, and evidentiary challenges, a consistent pattern emerges: the Act regulates risk procedurally but fails to constrain discriminatory outcomes substantively.

General-purpose AI models illustrate this systemic weakness. Their capacity to generate profiling outputs long before deployment escapes the AI Act's "high-risk" classification, leaving a gap between model-level and context-specific accountability. Similarly, the treatment of inferred sensitive data exposes a normative disconnect between Article 9 GDPR—which prohibits processing that *reveals* special-category data—and Article 10 AI Act, which merely governs the procedural quality of datasets. This allows sensitive inferences to persist under a veneer of compliance. Finally, the burden-of-proof problem highlights how the opacity of AI systems undermines victims' ability to establish discrimination, weakening the effectiveness of existing equality and data-protection frameworks.

To address these limitations and strengthen protection against discriminatory profiling, we suggest explicitly integrating the AI Act with the GDPR's substantive prohibitions and the EU equality directives' procedural guarantees. Equality-related risk should serve as a criterion for classifying systems as "systemic" or "high risk"; algorithmic inference of protected traits should trigger both Article 9 GDPR and Article 10 AI Act obligations; and evidentiary access should be enhanced through transparency and documentation duties at the model level.

In short, the technical design of AI systems cannot be separated from their legal regulation. For the AI Act to meet its fundamental-rights objectives, it should move beyond procedural safeguards toward a coherent, enforceable framework that directly addresses the technological sources of discriminatory profiling.

#### 4. Concluding Remarks

The coexistence of the AI Act, the GDPR, and the LED reflects the EU's ambition to regulate artificial intelligence through a multi-layered legal framework rooted in fundamental rights. Despite this, the analysis above shows that their combined operation remains fragmented and incomplete, particularly in addressing profiling and the generation of inferred sensitive data. The key challenge lies in the misalignment of regulatory logic: the AI Act's *ex ante* risk-management model operates at the system level, while the GDPR and LED guarantee *ex post* individual rights. Without integration mechanisms, these two regulatory dimensions fail to deliver coherent protection against automated discrimination.

Profiling based on inferred data exposes individuals to covert forms of categorisation that escape traditional notions of consent and transparency. The AI Act does not directly prohibit the generation of sensitive inferences, while the GDPR only reacts after harm occurs, assuming the individual can even detect it. The LED compounds the issue by allowing broad secrecy exceptions in law-enforcement profiling. The result is a fragmented regime that secures procedural compliance but not substantive equality.

For the EU's digital governance to meet its fundamental-rights commitments under Articles 8, 21, and 47 of the Charter, reforms and interpretative alignment are required. The following consolidated recommendations summarise the comment's proposed adjustments:

- First, AI governance should be structured so that it strengthens, rather than replaces, individual rights. System-level oversight under the AI Act should therefore include explicit safeguards ensuring that it complements, and does not dilute, core rights under the GDPR and the LED. In particular, the framework should guarantee meaningful human review, robust transparency, effective access to remedies, and judicial oversight, especially where automated profiling produces legal or similarly significant effects on individuals.
- Building on this foundation, regulatory gaps between the AI Act, the GDPR, the LED, and EU equality law should be formally closed. To prevent fragmentation and compliance arbitrage, the relevant regimes should be explicitly linked so that AI systems capable of inferring sensitive or otherwise protected characteristics automatically trigger high-risk classification and fundamental-rights obligations. Such

alignment would ensure consistent protection across both public and private sector deployments.

- To give practical effect to these aligned obligations, human oversight and accountability mechanisms must be made meaningful in practice. Human reviewers should possess genuine authority, relevant contextual expertise, and the power to override automated outputs. This should be supported by clear operational guidance from EU bodies and harmonised interpretative standards to avoid superficial or purely symbolic oversight.
- Complementing meaningful oversight, transparency and information rights should be strengthened, particularly in high-risk and law-enforcement contexts. Where operational needs justify temporary restrictions on transparency, minimum information “floors” and ex-post notification duties should be introduced. This would safeguard proportionality, ensure accountability, and preserve access to legal remedies once legitimate risks to investigations or security have subsided.
- Effective rights protection further depends on institutional coordination. Accordingly, cross-authority cooperation and shared enforcement mechanisms should be institutionalised. Formal structures—such as Memorandums of Understanding, joint task forces, coordinated reporting mechanisms, and mutual reliance on findings—should be established among the AI Office, Data Protection Authorities, market surveillance authorities, and equality bodies. Shared inspection tools and coordinated enforcement strategies would strengthen the response to discriminatory or unlawful AI practices.
- In parallel, systemic discrimination risks in AI design and deployment must be addressed more directly. Equality-based risk designations should be introduced for powerful AI models, alongside strengthened documentation and disclosure obligations. Discriminatory risk profiling—particularly where race, ethnicity, gender, or their proxies are used—should be expressly prohibited. Where credible harm is demonstrated, the burden of proof should shift to providers or deployers to show compliance with fundamental rights and equality standards.
- Finally, to reinforce legitimacy and long-term effectiveness, AI regulation should be made more accessible and democratically accountable. Technical thresholds and core concepts should be clarified in plain language, and periodic EU-level reporting on profiling practices and equality impacts should be mandated. Such measures would support informed policymaking, facilitate public scrutiny, and enhance trust in the evolving AI governance framework.