Athens Journal of Law - Volume 12, Issue 3, July 2026 – Pages 209-224

Ensuring Justice and Non-Discrimination in Automated Decision-Making: A Fundamental Rights Perspective

*By Doina Popescu Ljungholm**

I will start with a simple question: what happens to fundamental rights when an algorithm, not a person, decides whether you receive social benefits, a loan, or even the right to enter a country? This paper examines exactly that problem, in the context of automated decision-making (ADM) systems within the European Union. In short, I argue that the current legal safeguards primarily Article 22 of the GDPR and the new AI Act contain structural gaps that hit the most vulnerable the hardest: ethnic minorities, migrants, and persons with disabilities. Why? Because these systems are opaque, the right to be heard becomes an empty formality, and enforcement authorities lack both the technical training and the resources to verify how complex algorithms actually work. Drawing on CJEU and ECtHR case law, plus two landmark cases the Dutch SyRI system and the UK Home Office visa-streaming tool I argue for three institutional reforms: extending mandatory Fundamental Rights Impact Assessments to all high-risk commercial deployments, defining clearly what "meaningful human oversight" really means, and, not least, equipping supervisory authorities with the resources they need to conduct genuine technical audits.

Keywords: *artificial intelligence, automated decision-making, non-discrimination, fundamental rights, EU law*

Introduction

The deployment of automated decision-making (ADM) systems in domains of fundamental importance welfare allocation, credit assessment, border control, and employment raises a precise and pressing legal question: whether existing EU law provides adequate protection for the individuals subject to algorithmically generated determinations that affect their rights and interests. When a municipal algorithm classifies a resident as a probable welfare fraudster without her knowledge, or when a credit-scoring system declines an application without any human having reviewed the file, the issue at stake is not merely one of administrative inefficiency. It is one of accountability, opacity, and the structural capacity of legal frameworks to respond to a form of governance that is at once pervasive and resistant to conventional oversight. This paper is concerned with exactly that structural resistance.

Here is the central argument I want to make. The existing EU legal framework, for all its stated ambitions, is simply not built well enough to protect fundamental rights in the ADM context. Take the GDPR[1] and its prohibition of "solely automated" decisions under Article 22. That threshold dissolves the moment a deployer hires

*Associate Professor, National University of Science and Technology Politehnica Bucharest, Romania.

[1]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 (General Data Protection Regulation – GDPR), OJ L 119, 4.5.2016, p. 1.)

https://doi.org/10.30958/ajl.12-3-4 doi=10.30958/ajl.12-3-4

someone to rubber-stamp whatever the algorithm has already decided an empty formality. The AI Act[2] is more ambitious in its regulatory reach, but its heaviest obligations fall on system providers not on the commercial deployers whose systems most directly affect people's lives. What we end up with is a framework that speaks the language of human dignity and non-discrimination with considerable sophistication, yet offers the people most at risk legal protections that are, in practice, difficult to invoke and even harder to enforce.

The Charter of Fundamental Rights of the European Union[3] gives me the normative benchmark for the critique that follows. Articles 1, 21, and 47 human dignity, non-discrimination, effective remedy are primary EU law. Any secondary instrument that falls structurally short of those standards is, to the extent of that shortfall, legally deficient. And I argue that both the GDPR and the AI Act exhibit precisely that shortfall and that closing the gap requires institutional redesign, not just tinkering with legislation.

Who bears the costs of these deficiencies matters enormously. Ethnic minorities, migrants, persons with disabilities, and those living in poverty are overrepresented among those subject to algorithmic decision-making in welfare, employment, border control, and housing. In each of these domains, a wrong decision carries consequences of a completely different order of severity from, say, an incorrectly recommended playlist. Hildebrandt has argued convincingly, in my view that automated systems do not just passively reflect social realities; they actively generate them, embedding the biases of historical data into decisions about individual futures.[4] Rouvroy's concept of "algorithmic governmentality"[5] gives this dynamic its most precise theoretical formulation: what is produced is a mode of governance that operates as if it were making no normative choices and is therefore accountable for none of the choices it makes.

The accountability problem gets even worse when we add what Katyal calls the privatisation of consequential decisions,[6] and what Pasquale captured in that wonderful image of the "black box":[7] proprietary systems whose logic is inaccessible not only to the individuals they affect, but also to the courts asked to review their outputs, and to the regulators charged with overseeing their deployment. When you cannot know

---

[2]Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act), OJ L, 12.7.2024. The Act entered into force on 1 August 2024 and applies progressively until 2 August 2027.

[3]Charter of Fundamental Rights of the European Union, OJ C 326, 26.10.2012, p. 391. By virtue of Article 6(1) TEU, the Charter has the same legal value as the Treaties.

[4]Mireille Hildebrandt, Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology (Edward Elgar, 2015) 56–58

[5]Antoinette Rouvroy, 'The end(s) of critique: Data behaviourism versus due process' in Mireille Hildebrandt and Katja de Vries (eds), Privacy, Due Process and the Computational Turn (Routledge, 2013) 143–167. The concept of 'algorithmic governmentality' draws on Foucault's analysis of governmental rationality to describe a mode of rule that operates through the automated production of norms from data, bypassing discursive justification and individual subjectivity.

[6]Sonia Katyal, 'Private Accountability in the Age of Artificial Intelligence' (2019) 66 UCLA Law Review 54, 59–63.

[7]Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (Harvard University Press, 2015) 3–8.

the basis of a decision, the right to be heard becomes formal, not real – a procedural entitlement without substantive content.

There is also a specific discrimination problem that deserves separate attention. Classical anti-discrimination law is built on identifying a protected characteristic and showing differential treatment on that ground. Algorithmic systems frequently produce discriminatory outcomes through a different mechanism: proxy variables – postcode, browsing history, purchasing patterns that are neutral on their face but correlate statistically with race, disability, or social class. Gerards and Zuiderveen Borgesius have shown just how comprehensively this mechanism evades existing legal categories,[8] and Eubanks has documented the human consequences of this phenomenon with an ethnographic precision that legal scholarship rarely achieves.[9]

The Court of Justice of the European Union has not been idle, to be fair. The SCHUFA judgment of December 2023[10] settled a significant interpretive dispute by holding that automated credit-scoring falls within Article 22 of the GDPR where it functions as a determinative input to a third party's decision whatever the formal structure of that decision might be. The Dun & Bradstreet Austria ruling of February 2025[11] moved the transparency question forward, establishing that trade secrecy cannot serve as a blanket defence against a data subject's right to understand how an automated decision was reached. Both rulings matter. But each operates at the level of individual challenge – after the harm has occurred, in cases where the affected person had the resources, knowledge, and stamina to litigate. The ECtHR Grand Chamber in D.H. and Others v Czech Republic[12] recognised, in a different context, that discrimination of a structural kind demands structural responses. That recognition has not yet found its way into the operational design of EU AI governance.

What follows maps the legal framework, evaluates how adequate it is for vulnerable groups, and proposes three institutional reforms: extending mandatory Fundamental Rights Impact Assessments to commercial ADM deployments; defining meaningful human oversight in substantive operational terms; and resourcing supervisory authorities for genuine technical audit.

---

[8]Janneke Gerards and Frederik Zuiderveen Borgesius, 'Protected Grounds and the System of Non-Discrimination Law in the Context of Algorithmic Decision-Making and Artificial Intelligence' (2020) 29 Information & Communications Technology Law 303, 305–310.

[9]Virginia Eubanks, Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor (St. Martin's Press, 2018) 11–14.

[10]CJEU, Case C-634/21, OQ v Land Hessen (SCHUFA Holding AG), ECLI:EU:C:2023:957, Judgment of 7 December 2023. The Court held that an automated credit-score constitutes an 'automated individual decision' within Article 22(1) GDPR where third parties determinatively rely on it.

[11]CJEU, Case C-203/22, Dun & Bradstreet Austria GmbH v CK, ECLI:EU:C:2025:137, Judgment of 27 February 2025. The Court held that national law cannot categorically exclude access to explanations in favour of trade secrets; controllers must provide the 'procedure and principles actually applied' in an intelligible and accessible form.

[12]ECtHR, D.H. and Others v Czech Republic [GC], Application No. 57325/00, Judgment of 13 November 2007, paras 175–180. Indirect discrimination producing disproportionately adverse effects on a group falls within Article 14 ECHR; statistical evidence may establish a prima facie case

Literature Review

Scholarship on ADM and fundamental rights has grown into a genuinely interdisciplinary conversation – one that spans computer science, legal theory, political philosophy, and empirical social research. The conversation is productive but uneven. Technical findings on algorithmic harm accumulate faster than the normative frameworks needed to evaluate them. Doctrinal legal analysis of GDPR and AI Act provisions proceeds, with some exceptions, without adequate engagement with how deployed systems actually behave. This review identifies three bodies of work directly relevant to the argument developed below and closes with an account of the gaps that give this paper its reason for being.

The Right to Explanation: From Promise to Paradox

Perhaps the most consequential single contribution to the legal-technical literature on ADM is Wachter, Mittelstadt and Russell's 2018 paper.[13] What made it consequential was its willingness to say plainly what the GDPR does and does not provide: a right to meaningful information about the logic involved in automated decisions, certainly but not a right to an explanation in any sense that would allow an affected individual to understand, challenge, or replicate the decision. A controller satisfies Article 15(1)(h) by communicating the general factors that influenced an output, without disclosing how those factors were weighted or combined. The remedy the authors propose counterfactual explanations that describe the minimum change to the input that would have produced a different result is elegant precisely because it threads the needle between transparency and commercial confidentiality. It is also, as the CJEU's Dun & Bradstreet Austria ruling suggests, increasingly influential on judicial thinking, though without explicit acknowledgment.

The same authors' 2021 paper carries a harder message.[14] The major statistical fairness metrics demographic parity, equalised odds, individual fairness are mutually incompatible when base rates differ across groups. This is a mathematical result, but its legal implications are severe: any system optimised to satisfy one fairness definition necessarily violates another, and EU non-discrimination law gives us no guidance on which definition should prevail. The AI Act inherits this silence. Its high-risk system requirements mandate bias monitoring and testing, but leave open the question of what fairness metric the monitoring is designed to detect. Without a principled answer to that question, the regulatory obligation is formally present and substantively empty. Edwards and Veale put the point from a different angle:[15] demanding explanations of discriminatory decisions does not make those decisions less discriminatory. It relocates the problem from the system to the individual, who must now obtain,

[13]Sandra Wachter, Brent Mittelstadt and Chris Russell, 'Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR' (2018) 31(2) Harvard Journal of Law and Technology 841–887.

[14]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 105567.

[15]Lilian Edwards and Michael Veale, 'Slave to the Algorithm? Why a Right to an Explanation Is Probably Not the Remedy You Are Looking For' (2017) 16 Duke Law & Technology Review 18, 20–24.

interpret, and legally operationalise an explanation a burden that is, for the populations most affected, prohibitive in practice.

Algorithmic Discrimination: Mapping the Harm

Zuiderveen Borgesius's Council of Europe study remains the most comprehensive mapping of algorithmic discrimination risk across sectors.[16] Its core finding is so direct it is worth preserving in plain language: anti-discrimination law cannot, by itself, adequately address algorithmic bias, because the causal structure of that bias does not match the causal structure anti-discrimination law was built to address. Machine learning systems trained on historically discriminatory data acquire proxies – postcode, purchasing behaviour, web browsing patterns that reproduce discriminatory effects without replicating discriminatory intent or even using protected categories as inputs. The law, as it stands, regulates inputs; the harm occurs through outputs. Closing that gap requires different law, not better algorithms. The computer science tradition's foundational concept of individual fairness[17] similar individuals should be treated similarly offers a partial conceptual bridge, but the operationalisation of "similarity" involves normative judgments that cannot be read off from technical specifications and that AI Act provisions on bias testing currently leave entirely unresolved.

The Regulatory Literature: Advances and Structural Limits

Two recent papers in the Athens Journal of Law take up precisely the structural tensions that this paper analyses. Sarra's examination of the relationship between AI Act Article 14 and GDPR Article 22[18] identifies what may prove to be the most consequential unintended consequence of the new architecture. The AI Act's mandatory human oversight requirement was designed as a safeguard against fully automated decision-making. But "solely automated" is the threshold for Article 22 GDPR's protections. A formally inserted human reviewer even one whose involvement is perfunctory removes a decision from Article 22's scope.

The human oversight requirement, intended as protection, operates as an exemption. Sarra's proposed re-interpretation of Article 22 focusing on the substantive independence of human review from the algorithmic output, rather than on its formal presence is, in my assessment, the most practically workable solution currently on offer. Boura's regulatory analysis[19]identifies a different structural choice: the sandbox model that allows providers to test AI systems in real-world conditions with reduced compliance obligations places the cost of technological development on the individuals who interact with those systems – frequently in welfare, healthcare, and criminal justice contexts, exactly those settings where the individuals concerned are most vulnerable and least

[16]Frederik Zuiderveen Borgesius, Discrimination, Artificial Intelligence, and Algorithmic Decision-Making (Council of Europe, Directorate General of Democracy, 2018) 13–18.
[17]Cynthia Dwork and others, 'Fairness Through Awareness' (Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ACM, 2012) 214–226.
[18]Claudio Sarra, 'Artificial Intelligence in Decision-making: A Test of Consistency between the EU AI Act and the General Data Protection Regulation' (2025) 11(1) Athens Journal of Law 45–62.
[19]Marta Boura, 'The Digital Regulatory Framework through EU AI Act: The Regulatory Sandboxes' Approach' (2024) 10(3) Athens Journal of Law 385–398.

able to refuse. Convention 108[20] provides the international floor: its provisions on automated decision-making operate regardless of whether the decision is "solely" automated, offering a degree of protection that EU secondary legislation has not yet consistently reached.

Gaps and the Contribution of This Paper

Reading this literature as a whole, three gaps become apparent. The Charter of Fundamental Rights legally binding primary law, against which all EU secondary legislation may be measured and found wanting appears in the scholarship largely as background aspiration rather than operative legal standard. The category of "vulnerable groups" is invoked regularly but examined concretely in relatively few legal studies; the human reality behind the doctrinal categories remains underexplored. And the field has produced substantial critique of existing law alongside limited constructive synthesis: what would a rights-compliant ADM framework actually look like in institutional detail? The following sections attempt, if not to answer that question definitively, at least to make it less avoidable.

Methodology

This study employs a legal-doctrinal methodology, which is the standard approach in European Union fundamental rights research when the object of inquiry is the interpretation, coherence, and application of legal norms. The core of the analysis consists of a systematic examination of primary EU law (the Charter of Fundamental Rights), secondary legislation (GDPR and AI Act), and the interpretive case law of the CJEU and the ECtHR. To test how these legal provisions actually work in reality, I use two comparative case studies the Dutch SyRI welfare-fraud detection system and the UK Home Office visa-streaming tool selected because each reveals, with unusual documentary clarity, a specific mechanism through which legal gaps translate into concrete rights violations. The selection of these two cases requires a brief methodological note. Both SyRI and the UK streaming tool are exceptional in one important respect: each attracted litigation and investigative scrutiny that produced a documentary record sufficient to reconstruct the system's operation, its legal basis, and the response of public authorities when challenged. Most ADM deployments in high-risk domains do not generate comparable documentation, precisely because opacity is structurally incentivised. The analytical value of these cases lies not in their representativeness they are, in that sense, unusual but in what they make visible: SyRI exposes the consequences of permitting systemic opacity and placing the burden of challenge on individuals who are never informed that they have been assessed; the streaming tool demonstrates what occurs when existing anti-discrimination law is formally in force but no pre-deployment mechanism requires compliance to be verified. Together, they instantiate two distinct failure modes of the regulatory architecture analysed in Section 4, and they provide the empirical basis for the reforms proposed in Section

[20]Council of Europe, Convention 108+ on the Protection of Individuals with regard to Automatic Processing of Personal Data (modernised version, CETS No. 223, 2018), Articles 8 and 9.

6. The analysis is normative in orientation: it does not merely describe the current legal architecture but evaluates it against the benchmark of the Charter's guarantees of human dignity, non-discrimination, and effective remedy.

The EU Legal Framework: Architecture and Fault Lines

The EU legal framework governing ADM is layered: primary law establishes the normative standard, secondary legislation attempts to operationalise it, and judicial interpretation fills or tries to fill the gaps that drafting leaves behind. What I call "fault lines" are the points where these layers fail to connect: where the protection promised at one level evaporates at another, and where that evaporation consistently works to the benefit of the deploying institution rather than the affected individual. That consistency, I would argue, is not coincidental.

The GDPR: A Safeguard with Built-In Escape Routes

Article 22 of the GDPR[21] was a legislative first the first binding EU provision to address the risks of consequential automated decisions directly. Its limitations were visible from the start. The "solely automated" threshold is the most obvious: any formal human participation in the decision chain, however perfunctory, removes the case from Article 22's scope. The Article 29 Working Party's guidance[22] attempted to give "meaningful" human review some content, but stopped short of providing operational criteria that supervisory authorities could enforce against deployers unwilling to comply. A second escape route runs through Article 22(2)(a), which excepts decisions necessary for a contract a formulation broad enough to cover most credit, insurance, and employment ADM. The third problem is structural rather than definitional: Article 22 is reactive. It gives you a right to contest a decision already made, but imposes no obligation on controllers to assess, before deploying a system, whether it will produce discriminatory results. The harm arrives before the law is available to address it.

The AI Act: Real Advances, Conspicuous Exclusions

The AI Act represents a genuine step beyond the GDPR's data-protection-centred approach. Its risk-based architecture[23] subjects high-risk AI systems to pre-market conformity assessment, technical documentation requirements, and human oversight obligations. For certain deployers, Article 27 mandates a Fundamental Rights Impact

[21]GDPR, Articles 22(1)–(3). The 'solely automated' threshold has attracted persistent criticism precisely because it is circumvented so easily by nominal human participation. Article 22(2)(a) additionally excepts decisions necessary for a contract, which in practice covers most commercial ADM.

[22]Article 29 Working Party (now EDPB), Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679, WP251rev.01 (revised 6 February 2018). The WP29 insisted human review must be 'meaningful' rather than perfunctory, but stopped short of specifying criteria that supervisory authorities could operationalise against deployers.

[23]AI Act, Article 6 and Annex III. High-risk systems include those deployed in biometric identification, education, employment, essential services, law enforcement, migration, and the administration of justice.

Assessment before deployment.[24] Article 5's prohibitions on social scoring by public authorities and on systems that exploit group vulnerabilities[25] are categorical: no proportionality balancing is permitted. These are genuine advances. Their limitation lies in their scope. The FRIA obligation under Article 27 applies only to public-body deployers. Commercial deployers in credit, insurance, recruitment, and platform governance the sectors where algorithmic discrimination against vulnerable groups is best documented fall entirely outside its reach. The human oversight requirements of Article 14 govern system design rather than the quality of review actually practised. There is no mechanism to verify that the human reviewer inserted into a high-risk AI system's decision chain does anything more than approve the algorithmic output. The resulting regulation is, as I argue, most protective where protection is least needed.

The Charter and the Courts: Filling the Gaps Cautiously

Judicial interpretation has begun to give the Charter's guarantees operative content in the ADM context, though progress is sector-specific and reactive. The SCHUFA judgment[26] clarified the scope of Article 22 GDPR in a way the legislature had left ambiguous: an automated score that functions as the determinative input to a third party's decision is itself an automated decision for Article 22 purposes, regardless of how the subsequent decision is formally structured. The Dun & Bradstreet Austria ruling[27] addressed the transparency question: commercial secrecy is not a sufficient reason to deny a data subject access to a meaningful explanation of how an automated decision was reached. The Ligue des droits humains judgment[28]extended the analysis to border-control profiling, holding that systematic automated risk classification constitutes a Charter-level interference requiring necessity and proportionality justification. Each of these rulings closes one escape route. None of them, however, operates upstream of the harm. They provide legal recourse after a wrong decision has been made and in the rare case successfully challenged. The ECtHR's D.H. judgment[29] points toward a more demanding model: structural discrimination, demonstrated statistically, requires systemic institutional responses. That model has not yet been translated into the operative architecture of EU secondary law on AI.

[24]AI Act, Article 27(1). The FRIA obligation is confined to public-body deployers and private entities providing public services — a formulation that leaves commercial deployers in credit, insurance, recruitment, and platform governance outside its mandatory scope.

[25]AI Act, Article 5(1)(c) and (e). These prohibitions are absolute: no proportionality balancing is available. They cover AI systems that exploit vulnerabilities of specific groups and social scoring by public authorities that produces detrimental treatment.

[26]CJEU, Case C-634/21, SCHUFA (n 10), paras 50–63.

[27]CJEU, Case C-203/22, Dun & Bradstreet Austria (n 11), para 71.

[28]CJEU, Case C-817/19, Ligue des droits humains v Conseil des ministres, ECLI:EU:C:2022:491, Judgment of 21 June 2022. Systematic automated PNR profiling was held to constitute a serious interference with Articles 7 and 8 of the Charter, requiring strict necessity and proportionality even where no individual decision is formally automated.

[29]ECtHR, D.H. and Others v Czech Republic (n 12), para 175.

Case Studies: When Law Meets Algorithm

The two cases examined here were chosen for what they reveal rather than for what they represent. Both are exceptional each attracted litigation and public attention that most ADM deployments do not. What makes them analytically useful is that each exposes, in documentary detail, a specific failure mode of the regulatory architecture described in Section 4. SyRI shows the consequences of permitting opacity and placing the enforcement burden on individuals who do not even know they have been harmed. The UK streaming tool shows what happens when the law clearly prohibits discriminatory conduct but imposes no obligation to check for compliance before deployment.

SyRI: The Algorithm the State Refused to Explain

Between 2014 and 2020, the Dutch government ran a system called SyRI Systeem Risico Indicatie which drew on data from tax authorities, municipal housing records, employment registers, and immigration files to generate risk scores identifying individuals as potential welfare fraudsters. The system operated without the knowledge of those it processed. People who received a risk score were not notified; they had no opportunity to see the score, understand how it had been calculated, or challenge it before investigators arrived. They were, to use the language of administrative law, the objects of a consequential classification made entirely behind their backs.In February 2020, the District Court of The Hague ruled SyRI unlawful.[30] The court grounded its ruling in Article 8 ECHR – the right to private and family life rather than in the GDPR, which had been in force for two years by that point and whose enforcement architecture had produced no action in six years of SyRI's operation. The court's central finding was that SyRI legislation failed the Convention's "quality of law" requirement: it was insufficiently clear about which data combinations could justify a fraud risk conclusion. More strikingly, the Dutch state had refused to disclose how the algorithm worked, even to the court itself. Van Bekkum and Zuiderveen Borgesius note the judgment's limited immediate effect:[31] the Dutch government subsequently developed successor systems under different legal bases.

Rachovitsa and Johann identify the deeper wound: "intentional opacity" designing a system so that its operation cannot be externally verified transforms the right to an effective remedy from a substantive guarantee into a procedural formality.[32]Both the

[30]District Court of The Hague, NJCM et al. v The Dutch State, ECLI:NL:RBDHA:2020:1878, Judgment of 5 February 2020, paras 6.7–6.9. The state's refusal to disclose the algorithm's logic even to the court prevented judicial verification of discriminatory operation, which the court identified as independently problematic.

[31]Marvin van Bekkum and Frederik Zuiderveen Borgesius, 'Digital Welfare Fraud Detection and the Dutch SyRI Judgment' (2021) 23 European Journal of Social Security 1, 5–8. The authors argue persuasively that without a general doctrine of algorithmic accountability, individual court victories tend to produce system redesign rather than structural reform.

[32]Adamantia Rachovitsa and Niclas Johann, 'The Human Rights Implications of the Use of AI in the Digital Welfare State: Lessons Learned from the Dutch SyRI Case' (2022) 22(2) Human Rights Law Review ngac010. The authors' concept of 'intentional opacity' designing a system so that its operation cannot be externally verified is a valuable analytical contribution that this paper extends to the enforcement context.

ECHR and the GDPR, correctly applied, would have prohibited SyRI. The failure was one of enforcement, and it was systemic: it reflected the structural position of a regulatory architecture that places the burden of challenge on the person who does not know she has been wronged.

The UK Streaming Tool: What Existing Law Required – and Was Never Applied

The UK Home Office's visa-streaming algorithm ran from 2015 to 2020. Over those five years, it assessed virtually every visa application by assigning a traffic-light risk rating green, amber, red that determined the intensity of scrutiny an application received and, consequently, its probability of refusal. The Home Office admitted, when pressed, that nationality was a primary variable: applicants from countries the algorithm designated "suspect" received higher risk scores, longer processing times, and substantially higher refusal rates. This is direct discrimination on grounds of national origin. The Equality Act 2010, section 13, unambiguously prohibits it. The Public Sector Equality Duty under section 149 required the Home Office to have due regard to the need to eliminate such discrimination before adopting the system. GDPR Article 35, operative from May 2018, required a Data Protection Impact Assessment for any processing likely to produce high risks to individuals' rights and freedoms.

None of these obligations were discharged before deployment. When JCWI and Foxglove filed for judicial review in June 2020,[33] they were asserting rights that had been in force for between two and ten years. The Home Office discontinued the tool before a full hearing, committing to conduct the equality and data protection assessments that should have preceded the system's original deployment.[34] No binding legal determination of the tool's unlawfulness was ever made. This is the most important aspect of the case, and the one most consistently overlooked when it is cited as an example of algorithmic accountability in action. It is, equally, an example of a public authority deploying a discriminatory system in confident knowledge that no enforcement mechanism would require prior justification and being proved right. The AI Act's ex ante FRIA requirement, had it existed and applied, would have forced the question before the harm was done. It did not exist; it now does; it applies from August 2026; it does not apply to commercial deployers; it offers no retroactive protection to anyone processed by any algorithm before that date.

[33]JCWI and Foxglove v Secretary of State for the Home Department, judicial review filed June 2020. The claim alleged direct racial discrimination under the Equality Act 2010, s. 13, and breach of the Public Sector Equality Duty under s. 149.

[34]Home Office letter to JCWI, 3 August 2020, discontinuing the Streaming Tool 'pending a redesign of the process'. The commitment to conduct Equality Impact Assessments and Data Protection Impact Assessments standard obligations already operative under the Equality Act 2010 and the GDPR — was presented as a new undertaking, which speaks for itself.

Discussion: Three Reforms and Their Discontents

A pattern runs through the legal framework analysis and the case studies alike, and it is not ambiguous. The existing architecture fails in a consistent direction: it protects the deploying institution and exposes the individual whose rights are at stake. SyRI ran for six years – not because the law permitted it, as the court confirmed, but because enforcement mechanisms placed the burden of challenge on people who did not know they had been harmed. The streaming tool ran for five years encoding direct discrimination not because the law permitted it, but because no authority required an equality assessment before deployment. Article 22 GDPR is circumvented daily through nominal human insertion not by legislative intent, but because no one has defined what meaningful review actually requires. The word for a pattern of failure that consistently advantages the powerful over the vulnerable, and that is predictable from the design of the system producing it, is structural.

Three reforms are necessary. I offer them here as minimum conditions not a complete programme for closing the gap between formal compliance and substantive rights protection.

ADM Risk Mitigation for the Protection of Fundamental Rights: Existing and Prospective Mechanisms

Before setting out the three proposed institutional reforms, it is necessary to map the risk mitigation mechanisms already available under EU law and to identify where each one falls structurally short. This mapping provides the analytical basis for the reforms that follow and establishes that the proposals are not alternatives to existing instruments but responses to their demonstrated inadequacies.

The primary existing mitigation mechanism is the Data Protection Impact Assessment (DPIA) required under GDPR Article 35 for processing likely to result in high risk to individuals' rights and freedoms. The DPIA obligation is, in principle, a pre-deployment instrument: it requires controllers to identify and assess risks before a system goes live. In practice, as the UK streaming tool case illustrates, it is routinely omitted without triggering enforcement action. Its limitation is structural: the DPIA framework is internally assessed, controller-conducted, and subject to no independent technical verification requirement. A controller may satisfy Article 35 on paper without any substantive engagement with the algorithmic processes at issue.

The AI Act introduces two significant mitigation mechanisms for high-risk systems: the conformity assessment under Article 43, which requires technical documentation, risk management procedures, and bias testing before market placement; and the Fundamental Rights Impact Assessment (FRIA) under Article 27, which requires deployers to identify the foreseeable impact of a high-risk AI system on fundamental rights. Both mechanisms represent genuine advances. The conformity assessment, in particular, establishes ex ante obligations that did not exist under the GDPR. Their limitation, as noted in Section 4.2 above, is one of scope: the FRIA is mandatory only for public-body deployers, and conformity assessments in most high-risk categories are conducted by providers rather than independently verified by third parties.

A third mechanism operates at the level of individual redress: the right to contest automated decisions under GDPR Article 22, and the right to an effective remedy under Charter Article 47. These are reactive instruments they operate after a harm has occurred and require the affected individual to initiate proceedings. As the SyRI case demonstrates, where a system is designed to be opaque and individuals are not informed that they have been assessed, these rights are practically unavailable to precisely those who most need them. Convention 108+ provides a floor that operates independently of the "solely automated" threshold, but its enforcement mechanism, at the level of national supervisory authorities, reproduces the resource and capacity constraints that limit GDPR enforcement. It is against this background of mechanisms that are real but structurally insufficient — that the three reforms proposed below are advanced.

Three Proposed Reforms

First reform: extend the mandatory Fundamental Rights Impact Assessment (FRIA) obligation to all high-risk commercial ADM deployments, with standardised methodology, independent verification, and mandatory public disclosure of results. The predictable objection is cost and innovation burden. It should be answered directly: pharmaceutical products, financial instruments, and major infrastructure all require pre-deployment risk assessment as a condition of market access.

The AI systems that determine access to credit, housing, employment, and welfare affect people's lives as consequentially as those products. The question worth asking is why, until 2024, they did not. It is, however, important to be clear-eyed about the institutional and political constraints that a reform of this scope would face. Legislative inertia within EU co-decision procedures, combined with the divergent interests of Member States whose domestic AI industries would bear the compliance costs, creates a structural tendency toward diluted obligations and extended phase-in periods. Industry lobbying by technology firms and financial sector actors who have invested substantially in ADM infrastructure and whose competitive position would be affected by mandatory pre-deployment assessment represents a further source of resistance that should be anticipated rather than discounted. The appropriate response is not to soften the proposal but to design the implementation architecture in a way that distributes costs equitably, allows for proportionate compliance frameworks for smaller operators, and builds in review mechanisms that allow the methodology to be refined as technical audit practice matures.

Second reform: address the human oversight problem Sarra identifies.[35] The AI Act requires that high-risk systems be designed to permit effective oversight. It does not define what effective oversight requires of the human reviewer. A system that presents the reviewer with a dashboard showing an algorithmic recommendation, without providing access to the underlying data, without institutional protection for a reviewer who chooses to override the recommendation, and without a requirement that the substance of the review not merely its outcome be recorded, satisfies the Act's letter and defeats its purpose. A minimum operational standard for meaningful review

[35]Sarra (n 18) 58–60. Sarra's proposal — that Article 22 GDPR be re-read to focus on the substantive independence of human review from the algorithmic output, rather than on formal participation — is one of the more practically workable suggestions in recent literature on this point.

is achievable; it requires political will to impose it on deploying organisations that have strong commercial reasons to prefer the current arrangement.

Third reform: resource supervisory authorities for genuine technical audit. Colonna documented in 2019 that no EU data protection authority had successfully audited a complex machine-learning system and imposed sanctions for Article 22 violations.[36] The AI Act creates new market surveillance authorities but does not ensure they have the technical staff, the financial resources, or – crucially – the institutional independence from government needed to audit AI systems deployed by public authorities. An enforcement body that cannot interrogate a training dataset, reproduce an algorithmic decision, or identify the point in the data pipeline where discriminatory proxies were learned cannot hold anyone accountable. Legislation without enforcement is not law; it is aspiration with a statutory citation.

A further structural issue in this third reform concerns the relationship between the AI Act's newly established market surveillance authorities and the national data protection authorities (DPAs) that already exercise enforcement jurisdiction under the GDPR. The AI Act does not resolve this relationship with adequate clarity. Where an AI system constitutes both a high-risk AI system under Annex III of the AI Act and involves automated processing of personal data under the GDPR which will be the case for most high-risk ADM deployments in welfare, credit, and border control both the market surveillance authority and the competent DPA have arguable jurisdiction.

The potential for overlap is significant: a single algorithmic deployment could in principle be subject to conformity assessment oversight by one authority and Article 22 GDPR enforcement by another, with neither authority having comprehensive visibility over the system's operation.

The reform proposed here therefore encompasses not only the resourcing question but also the coordination architecture: a statutory duty of cooperation between market surveillance authorities and DPAs, with clearly assigned lead competence for different enforcement functions, is a necessary complement to any resourcing uplift.Without it, the fragmentation of oversight reproduces at the institutional level the same accountability gaps that characterise the regulatory framework at the legislative level.

Conclusions

The cases and doctrinal analysis presented in this paper converge on a finding that is both specific and systemic. The Dutch welfare claimant whose risk score was generated by a system she was never informed of, the visa applicant whose file was pre-classified as high-risk on grounds of nationality before any official had reviewed it, and the loan applicant whose refusal was produced by a model whose logic was withheld as a commercial secret these are not isolated administrative failures. They are specific instances of a structural relationship between algorithmic power and

[36]Liane Colonna, 'Automated Decision-Making, Profiling, and the GDPR' (2019) 35 Computer Law & Security Review 397, 402–404. Writing before the AI Act's adoption, Colonna was already identifying enforcement capacity as the critical weakness; the subsequent years have done little to address her diagnosis.

legal accountability that the current EU framework has not adequately resolved. What the analysis reveals is not that the law is absent, but that the mechanisms through which it is supposed to operate are systematically inadequate to the task assigned to them.

The failure this paper has traced is architectural rather than merely legislative. Both the GDPR and the AI Act express genuine commitments to human dignity and non-discrimination; the problem lies in the gap between those commitments and the operative mechanisms through which they are supposed to be realised. The human-in-the-loop loophole converts a substantive safeguard into a procedural formality. The commercial exclusion from the FRIA obligation exempts from the most demanding requirements precisely the deployments that are most consequential for vulnerable people. The enforcement gap between what the law requires and what supervisory authorities can verify is one that deployers rationally exploit, because the cost of non-compliance falls on individuals rather than institutions.

The Charter of Fundamental Rights demands non-discrimination and human dignity with the same legal force as the Treaties. A regulatory architecture that cannot deliver those guarantees in practice is, to that extent, in breach of primary EU law – not as a matter of political rhetoric but as a matter of legal analysis. The task ahead is institutional: building the mechanisms through which commitments already made are actually kept. It would, however, be analytically incomplete to advocate for these reforms without acknowledging the practical challenges their implementation would face. The extension of mandatory FRIA obligations to commercial deployers will encounter resistance not only from industry actors but from Member States whose administrations benefit from the current permissive architecture. The definition of meaningful human oversight in operationally enforceable terms requires the development of audit standards and professional capacity that does not yet exist at scale across EU supervisory authorities.

The resourcing and coordination reforms proposed for supervisory authorities will require sustained budgetary commitment from Member States that have, historically, been reluctant to fund data protection enforcement at the level the GDPR's ambitions require. None of these challenges is insuperable, but none is trivial. The implementation of each proposed reform will require not only legislative amendment but sustained political will, administrative investment, and the development of technical expertise within public institutions that currently lack it. Acknowledging these constraints is not a reason to lower the normative ambition of the proposals; it is a reason to design the transition architecture with realism and to build in the review mechanisms that allow reforms to be calibrated as implementation experience accumulates. That task is feasible, and it is urgent. For those whose rights have already been affected by the inadequacies the current framework has failed to address, the case for prompt and serious action needs no further elaboration.

References

Academic Works

Boura, M. (2024). The Digital Regulatory Framework through EU AI Act: The Regulatory Sandboxes' Approach. *Athens Journal of Law* 10(3): 385–398.

Colonna, L. (2019). Automated Decision-Making, Profiling, and the GDPR. *Computer Law & Security Review* 35(4): 397–410.

Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). F*airness Through Awareness.* In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. New York: ACM, 214–226.

Edwards, L., & Veale, M. (2017). Slave to the Algorithm? Why a Right to an Explanation Is Probably Not the Remedy You Are Looking For. *Duke Law & Technology Review* 16: 18–84.

Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.

Gerards, J., & Zuiderveen Borgesius, F. (2020). Protected Grounds and the System of Non-Discrimination Law in the Context of Algorithmic Decision-Making and Artificial Intelligence. *Information & Communications Technology Law* 29(3): 303–333.

Hildebrandt, M. (2015). *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology*. Cheltenham: Edward Elgar.

Katyal, S. (2019). Private Accountability in the Age of Artificial Intelligence. *UCLA Law Review* 66: 54–141.

Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.

Rachovitsa, A., & Johann, N. (2022). The Human Rights Implications of the Use of AI in the Digital Welfare State: Lessons Learned from the Dutch SyRI Case. *Human Rights Law Review* 22(2): ngac010.

Rouvroy, A. (2013). The end(s) of critique: Data behaviourism versus due process. In Hildebrandt M and de Vries K (eds) *Privacy, Due Process and the Computational Turn*. London: Routledge, 143–167.

Sarra, C. (2025). Artificial Intelligence in Decision-making: A Test of Consistency between the EU AI Act and the General Data Protection Regulation. *Athens Journal of Law* 11(1): 45–62.

Van Bekkum, M., & Zuiderveen Borgesius, F. (2021). Digital Welfare Fraud Detection and the Dutch SyRI Judgment. *European Journal of Social Security* 23(1): 1–18.

Wachter, S., Mittelstadt, B., & Russell, C. (2018). Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR. *Harvard Journal of Law and Technology* 31(2): 841–887.

Wachter, S., Mittelstadt, B., & Russell, C. (2021). Why Fairness Cannot Be Automated: Bridging the Gap between EU Non-Discrimination Law and AI. *Computer Law & Security Review* 41: 105567.

Zuiderveen Borgesius, F. (2018). *Discrimination, Artificial Intelligence, and Algorithmic Decision-Making*. Strasbourg: Council of Europe, Directorate General of Democracy.

Legislation and Legal Instruments

Article 29 Working Party (2018) Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679, WP251rev.01 (revised 6 February 2018). Brussels: European Commission.

Charter of Fundamental Rights of the European Union (2012) OJ C 326, 26.10.2012, p. 391.
Council of Europe (2018) Convention 108+ for the Protection of Individuals with regard to Automatic Processing of Personal Data (modernised version, CETS No. 223). Strasbourg: Council of Europe.
European Parliament and Council (2016) Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation — GDPR). OJ L 119, 4.5.2016, p. 1.
European Parliament and Council (2024) Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). OJ L, 12.7.2024.
United Kingdom (2010) Equality Act 2010, c. 15. London: Her Majesty's Stationery Office.

Case Law

Court of Justice of the European Union, Case C-817/19, *Ligue des droits humains v Conseil des ministres*, ECLI:EU:C:2022:491, Judgment of 21 June 2022.
Court of Justice of the European Union, Case C-634/21, *OQ v Land Hessen (SCHUFA Holding AG)*, ECLI:EU:C:2023:957, Judgment of 7 December 2023.
Court of Justice of the European Union, Case C-203/22, *Dun & Bradstreet Austria GmbH v CK*, ECLI:EU:C:2025:137, Judgment of 27 February 2025.
European Court of Human Rights, *D.H. and Others v Czech Republic* [GC], Application No. 57325/00, Judgment of 13 November 2007. Reports of Judgments and Decisions 2007-IV.
Netherlands, District Court of The Hague, *NJCM et al. v The Dutch State*, ECLI:NL:RBDHA: 2020:1878, Judgment of 5 February 2020.
United Kingdom, High Court of Justice, *JCWI and Foxglove v Secretary of State for the Home Department*, judicial review filed June 2020; discontinued by consent August 2020.