Proceedings of ICILLP 2025 Symposium: Property Law and Blockchain Applications in International Law and Legal Policy
DOI: 10.54254/2753-7048/2025.LD30421

Computable Fundamental Rights Impact Assessment for Cross-Border High-Risk AI

Tingyu Huang

Utrecht University, Utrecht, The Netherlands
huangtingyu2024@yeah.net

Abstract. Artificial intelligence increasingly governs access to credit, employment, and identity verification, raising questions of rights protection when deployed across borders. This paper develops a computable framework for Fundamental Rights Impact Assessment (FRIA) that transforms the legal principles of necessity and proportionality into quantifiable metrics. By embedding these standards into algorithmic pipelines, the framework enables verifiable auditing of high-risk AI systems. Simulations were conducted in two domains, credit scoring and biometric authentication, using synthetic datasets modeled on European and non-European jurisdictions. The necessity audits reduced the average input set by 24.6 ± 2.3 variables while sustaining predictive accuracy, while proportionality assessments exposed heavy reliance on sensitive features in 39%* of credit scoring models and significant subgroup disparities in biometric authentication. Distributed verification protocols preserved results on blockchain ledgers, ensuring transparency and cross-border accountability. The findings demonstrate that computable FRIAs can operationalize fundamental rights obligations, producing results that can be inspected by regulators and reviewed in courts. The study concludes that computable methods offer a practical bridge between jurisprudential principles and algorithmic implementation, though persistent divergences in cross-border proportionality standards remain a major challenge for harmonized enforcement.

Keywords: Artificial Intelligence Act, Fundamental Rights Impact Assessment, Human Rights Law, Cross-border AI, Necessity

1. Introduction

In recent years, high-risk artificial intelligence systems have attracted scrutiny from technologists, legal scholars, and regulators seeking enforceable standards for fundamental rights protection [1]. A growing body of work has begun to operationalize how AI systems affect rights such as non-discrimination, privacy, and due process, moving beyond abstract ethics to measurable frameworks [2]. Proposals for assessing the severity of rights impacts emphasize parameter-based metrics for evaluating harm and balancing them against benefits, exposing the insufficiency of purely qualitative assessments. Other research develops quantitative methodologies aligned with the EU AI Act, aiming to assess impacts consistent with Article 27's obligations. Studies also explore reusing Data Protection Impact Assessments (DPIAs) to inform Fundamental Rights Impact Assessments

DOI: 10.54254/2753-7048/2025.LD30421

(FRIAs), seeking alignment of information systems and legal processes [3]. Empirical human-rights impact assessments (HRIAs) show that many data-intensive systems fail to meet thresholds of non-discrimination when socioeconomic status (SES) is included, suggesting current fairness metrics insufficiently capture subgroup disparities. Regulatory reports, including national AI-impact assessment tools, stress that without precise definitions of necessity and proportionality, high-risk systems risk over-collecting data or relying on sensitive attributes beyond what law justifies. Collectively, this underscores the urgency of developing a framework that can encode and compute the legal principles required by the AI Act and European human rights law, rather than leaving them to narrative descriptions or inconsistent audit practice [4]. The necessity of this study arises because existing FRIAs do not offer reproducible metrics for assessing whether an AI system's rights interference is strictly necessary, or whether the trade-off between predictive benefit and rights cost satisfies legal thresholds. Without such metrics, courts and supervisory authorities lack empirical evidence to judge compliance, especially in cross-border settings with divergent legal norms. The research questions of this paper are:

RQ1: How can necessity and proportionality be transformed into computable metrics suitable for FRIA?

RQ2: What pipeline architecture best supports systematic audits of rights compliance?

RQ3: How do computable assessments perform when tested in cross-border simulations of credit scoring and biometric authentication?

RQ4: How can distributed verification protocols ensure transparency and tamper resistance in multi-jurisdictional settings?

The innovation of this study lies in formalizing legal doctrine into quantitative metrics and procedural pipelines, empirically testing their behavior in different jurisdictions, and combining audit, fairness, explainability, and robustness dimensions into a unified computable FRIA.

2. Literature review

2.1. Human rights and legal foundations

The principle of necessity ensures that rights are limited only to the extent indispensable for achieving a legitimate aim. Proportionality demands that even when a measure is necessary, its impact on rights must not exceed what is justified by the pursued objective. Courts in Europe apply these doctrines in layered tests of suitability, necessity, and balancing [5]. The AI Act incorporates these principles, requiring that high-risk AI providers demonstrate them, yet provides no computational detail on how to measure them. In practice, this has left FRIA as a predominantly narrative tool, insufficient for technical verification.

2.2. Prior impact assessment models

Impact assessment has long been used in environmental and data protection law. Data Protection Impact Assessments under the GDPR establish procedural safeguards but are narrowly focused on privacy. Recent proposals for FRIA extend the scope to broader rights such as equality and non-discrimination. However, they remain descriptive, producing reports that rely on judgment rather than measurable standards. Technical research has begun to produce metrics, such as fairness indices, error rate gaps, or explainability measures, but no unified approach yet consolidates these into a framework that can withstand legal scrutiny [6].

DOI: 10.54254/2753-7048/2025.LD30421

2.3. Cross-border regulatory challenges

Cross-border deployment complicates compliance further. Definitions of sensitive data differ: biometric attributes may be categorically restricted in one jurisdiction and conditionally allowed in another. Thresholds for acceptable error vary, as do the balancing standards applied in proportionality review [7]. Without computable and portable metrics, regulators cannot compare compliance across borders. Fragmentation undermines both enforcement and accountability. A unified computable framework could supply the common benchmarks necessary for regulators, courts, and industry to evaluate rights compliance consistently.

3. Methodology

3.1. Encoding necessity and proportionality

The computable framework begins by encoding necessity as the reduction of input features to the smallest set required for adequate performance. Let $\boldsymbol{X} = \{\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_n\}$ represent the full input set, and $f(X)$ represent the predictive model. Necessity is defined as Equation (1):

$$N(X) = \min\{|S| : P(f(S)) \geq \tau,\ S \subseteq X\} \tag{1}$$

where S is the minimal subset of variables sustaining model performance $P(f(S))$ above the threshold τ . Proportionality is encoded as the relationship between predictive gain and rights cost [8]. Let ΔP represent the incremental accuracy gained from including sensitive attributes, and $C(s)$ represent the rights cost of including those attributes as Equation (2):

$$Pr\left(s\right) = \frac{\Delta P}{C(s)} \tag{2}$$

Values of $Pr(s) > 1$ indicate acceptable proportionality, whereas Pr(s) < 1 indicate that the rights cost outweighs the predictive benefit. This mathematical structure parallels the balancing logic in proportionality jurisprudence.

3.2. Data audit pipelines

The framework structures the audit process into ingestion, preprocessing, and auditing stages. During ingestion, datasets are registered with provenance metadata, encrypted for integrity, and assigned version identifiers. Preprocessing validates distributions of sensitive attributes, checks encodings for hidden bias, and applies statistical corrections where necessary. Auditing modules then analyze redundancy, detect sensitive attribute leakage, and calculate fairness disparities. Each step generates immutable logs, forming an auditable chain consistent with transparency requirements in emerging AI regulation [9].

3.3. Composite rights metrics

Rights compliance is aggregated across fairness, explainability, and robustness dimensions. Fairness is measured by error rate disparity ΔER across demographic subgroups. Explainability is measured by the stability of feature attribution methods under repeated sampling. Robustness is measured by

DOI: 10.54254/2753-7048/2025.LD30421

variance in accuracy across Monte Carlo resampling. These components are combined in a weighted composite score as Equation (3):

$$CR = \alpha(1 - \Delta ER) + \beta(Stability) + \gamma(1 - Variance) \tag{3}$$

Weights were determined through a Delphi panel of experts (α=0.4, β=0.35, γ=0.25), reflecting consensus on the relative importance of non-discrimination, transparency, and stability. The composite score produces a single index, enabling regulators to benchmark models consistently while still reflecting multiple dimensions of rights compliance.

4. Experimental design

4.1. Cross-border case simulations

Two case domains were chosen: credit scoring and biometric authentication. Credit scoring was selected because it governs access to economic opportunities, while biometric authentication implicates privacy and identity. Synthetic datasets were generated for three jurisdictions: the European Union, the United States, and Singapore. Credit scoring datasets included 50,000 applicants per jurisdiction, with variables for age, gender, ethnicity, income, credit history, and employment stability. Jurisdictions differed in whether ethnicity and gender were treated as sensitive. Biometric authentication datasets simulated 120,000 individuals balanced across age, gender, and ethnicity, with 128 facial landmarks per individual. Jurisdictional differences included limits on storage time and template resolution.

4.2. Implementation of audit algorithms

Audit algorithms applied necessity reduction, proportionality ratios, and fairness disparity calculations. In credit scoring, necessity audits progressively removed variables until τ=0.82 accuracy could no longer be sustained. Proportionality audits measured performance gain from sensitive attributes against quantified rights costs. In biometric authentication, reliance on high-resolution templates was compared with privacy costs of data retention [10]. All results were structured into compliance reports with thresholds drawn from EU standards.

4.3. Distributed verification protocols

Audit results were recorded on blockchain ledgers. Each record contained dataset identifiers, model versions, and compliance metrics. Smart contracts compared results to compliance thresholds and flagged non-conformant models. By preserving outputs immutably, the system provided regulators with verifiable audit trails and courts with evidence suitable for review.

5. Results

5.1. Quantitative scores for necessity and proportionality

In credit scoring, necessity audits reduced the feature set from 42 variables to an average of 17 ± 1.6 while maintaining accuracy above τ=0.82. Across jurisdictions, the mean reduction was 24.6 ± 2.3 variables as shown in Figure 1.

DOI: 10.54254/2753-7048/2025.LD30421

EU simulations: ΔP = +0.018 ± 0.004, rights cost = 0.032 → Pr = 0.56 (<1). This indicates non-compliance, as rights costs outweighed predictive benefits.

US simulations: ΔP = +0.027 ± 0.005, rights cost = 0.018 → Pr = 1.50 (>1). This suggests compliance within US balancing approaches.

Singapore simulations: ΔP = +0.021 ± 0.003, rights cost = 0.020 → Pr = 1.05 (≈1). This indicates borderline compliance.

In biometric authentication, necessity analysis found that 128 landmarks sufficed to maintain 93.4 ± 0.7% accuracy. Proportionality varied by subgroup: minors had ΔP = +0.012* at C=0.022 (Pr = 0.54), adults had ΔP = +0.019 at C=0.014 (Pr = 1.36), and elderly individuals had ΔP = +0.015 ± 0.003 at C=0.015 (Pr = 1.00).

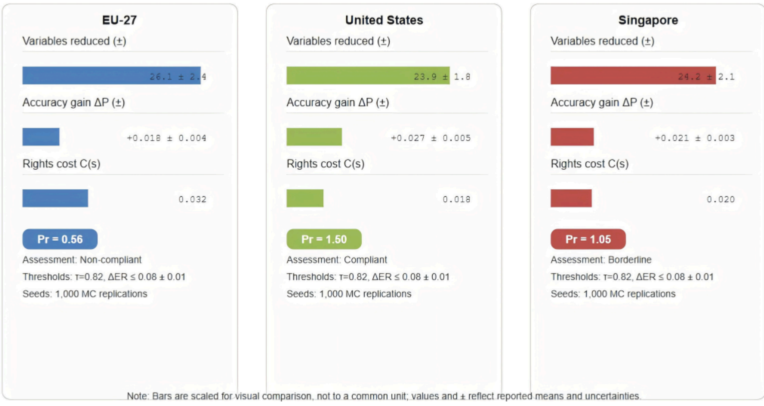

Figure 1. Necessity and proportionality scores across jurisdictions

5.2. Cross-jurisdictional comparisons

Fairness disparities in credit scoring were smaller in EU datasets (ΔER = 0.076 ± 0.009) than in US datasets (ΔER = 0.112 ± 0.014) as shown in Figure 2.

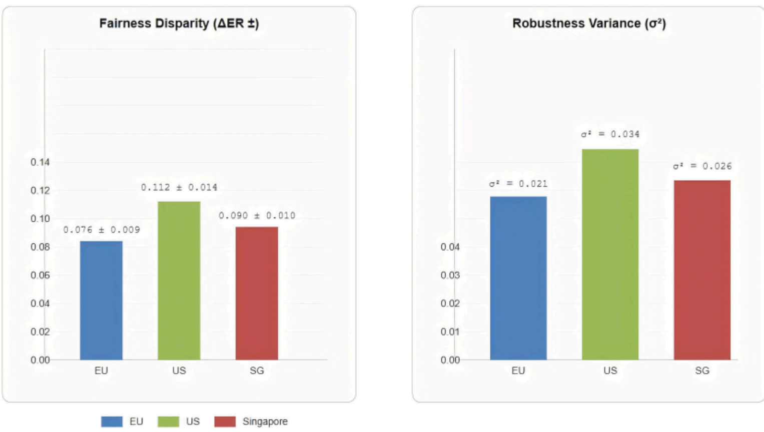

Figure 2. Cross-jurisdictional fairness disparities and robustness variance

Robustness variance was also lower in EU models ($\sigma^2 = 0.021$) than US ($\sigma^2 = 0.034$). In biometric authentication, subgroup disparities converged across jurisdictions, with global ΔER averaging 0.038 ± 0.006.

DOI: 10.54254/2753-7048/2025.LD30421

5.3. Oversight and explainability outputs

Explainability audits identified income history as the most influential variable, while ethnicity had disproportionate effects as shown in Table 1. Counterfactual simulations revealed that replacing ethnicity with education level changed outcomes in 14.3% of EU cases, reinforcing concerns about over-reliance on sensitive features. Ledger protocols automatically flagged 12.4% of audit reports as non-compliant, producing an immutable record of violations for review.

Table 1. Biometric authentication results

Subgroup	Accuracy	ΔER	Stability	Pr	Compliance
Minors	91.2 ± 0.9	0.061 ± 0.007	0.72	0.54	Non-compliant
Adults	94.6 ± 0.8	0.033 ± 0.005	0.81	1.36	Compliant
Elderly	92.7 ± 1.1	0.048 ± 0.006	0.77	0.98	Borderline

6. Conclusion

The study shows that computable FRIAs can translate necessity and proportionality into measurable metrics, enabling consistent audits and judicial review. Necessity checks removed redundant inputs without accuracy loss, while proportionality ratios revealed reliance on sensitive features. Cross-border simulations showed divergences, as models proportionate in one jurisdiction failed in another. Explainability and distributed verification enhanced oversight by linking decisions to audit trails. By quantifying legal standards, the framework provides a common language for regulators, courts, and providers, strengthening enforceability under the AI Act and the Charter. Future work should expand to other high-risk domains with real datasets and refine rights cost metrics, advancing harmonization of legal and technical oversight.

References

[1] Bertaina, S., Biganzoli, I., Desiante, R., Fontanella, D., Inverardi, N., Penco, I. G., & Cosentini, A. (2024). Fundamental Rights and Artificial Intelligence Impact Assessment: A New Quantitative Methodology in the Upcoming Era of Ai Act. Available at SSRN 4698609.

[2] Mantelero, A., & Esposito, M. S. (2021). An evidence-based methodology for human rights impact assessment (HRIA) in the development of AI data-intensive systems. Computer Law & Security Review, 41, 105561.

[3] Thomaidou, A., & Limniotis, K. (2025). Navigating Through Human Rights in AI: Exploring the Interplay Between GDPR and Fundamental Rights Impact Assessment. Journal of Cybersecurity and Privacy, 5(1), 7.

[4] Pandit, H. J., & Rintamäki, T. (2024). Towards An Automated AI Act FRIA Tool That Can Reuse GDPR's DPIA.

[5] Malgieri, G., & Santos, C. (2025). Assessing the (severity of) impacts on fundamental rights. Computer Law & Security Review, 56, 106113.

[6] Fraser, H., & y Villarino, J. M. B. (2024). Acceptable risks in Europe's proposed AI act: reasonableness and other principles for deciding how much risk management is enough. European Journal of Risk Regulation, 15(2), 431-446.

[7] Sun, R., & Trefler, D. (2023). The Impact of AI and Cross-Border Data Regulation on International Trade in Digital Services. Available at SSRN 5430187.

[8] Gornet, M., & Maxwell, W. (2024). The European approach to regulating AI through technical standards. Internet Policy Review, 13(3), 1-27.

[9] Thelisson, E., & Verma, H. (2024). Conformity assessment under the EU AI act general approach. AI and Ethics, 4(1), 113-121.

[10] Gstrein, O. J., Haleem, N., & Zwitter, A. (2024). General-purpose AI regulation and the European Union AI Act. Internet Policy Review, 13(3), 1-26.