AI Governance and Society Journal
ISSN: e-ISSN: 3117-6097 / p-ISSN: 3117-6100
https://galaxiauniverse.com/index.php/AIGSJ/issue/archive

Policy-Aware Computing Models for Ethical AI Compliance

Elena Muller[1]

[1] Associate Professor, Department of Computer Science, Central European Tech University, Vienna, Austria. Email: elena.muller27@gmail.com | ORCID: 3952-9915-5754-1597

ABSTRACT

Policy-aware computing -- the integration of formal policy representations into the computational architecture of AI systems such that compliance with ethical and regulatory obligations is enforced at the system level rather than dependent on human procedural adherence -- provides a promising technical paradigm for automating AI compliance at scale. As the EU AI Act (2024) and GDPR (2016) impose increasingly specific technical obligations on AI systems, the translation of regulatory text into machine-enforceable policy representations is a critical and underexplored engineering challenge. This paper proposes the Policy-Aware AI Compliance (PAAC) framework, a computational architecture for embedding EU AI Act and GDPR policy constraints into AI system operation through three integrated components: a Policy Formalisation Engine (PFE) that translates regulatory obligations into formal constraint specifications; a Runtime Policy Enforcement Layer (RPEL) that evaluates AI system actions against formalised constraints at inference time; and a Policy Compliance Audit System (PCAS) that maintains a cryptographically verified audit trail of all policy evaluations and enforcement decisions. The PAAC framework is evaluated through implementation on four production AI systems spanning healthcare, financial services, public sector, and legal AI domains, and through an expert assessment of 30 AI compliance and policy engineering specialists. Implementation results demonstrate 99.7% policy constraint coverage across 847 formalised EU AI Act and GDPR provisions, sub-100ms enforcement latency at inference, and a 67.4% reduction in compliance incidents over 12 months compared to manual compliance procedures. Expert consensus (Kendall $W = 0.76$, $p < 0.001$) confirms strong agreement on PAAC component importance. The study contributes the PAAC specification, a Policy Formalisation Language (PFL) for AI regulatory constraints, and empirical evidence that policy-aware computing substantially reduces AI compliance incident rates in production deployments.

Keywords: policy-aware computing; AI compliance; regulatory formalisation; runtime enforcement; EU AI Act; GDPR; responsible AI; compliance automation

Citation: Muller [2025]. Policy-Aware Computing Models for Ethical AI Compliance. DOI: https://doi.org/10.5281/zenodo.19185100

Article Information: *Received:* 30 August 2025 *Accepted:* 02 November 2025 *Published:* 22 December 2025

Research Article: *Research Article*

1. Introduction

1.1 The AI Compliance Automation Imperative

The governance of AI systems in regulated sectors is increasingly a compliance engineering problem as much as a policy design problem. The EU AI Act (2024) imposes over 200 specific technical and organisational obligations on high-risk AI system developers, deployers, and importers; the GDPR (2016) adds over 100 further data protection obligations applicable to AI systems processing personal data. The combined compliance surface of these instruments for a healthcare AI developer, for example, encompasses obligations spanning data governance (EU AI Act Article 10, GDPR Articles 5-9), system transparency (EU AI Act Article 13), human oversight (Article 14), risk management (Article 9), accuracy and robustness (Article 15), and post-market monitoring (Article 72). Manual compliance with this obligation landscape -- through policy documentation, procedural controls, and periodic audits -- is increasingly inadequate for three reasons. First, the volume and specificity of obligations exceed human attention capacity in fast-paced AI development cycles. Second, procedural compliance depends on individual adherence and is not verifiable at the granularity of individual AI inferences. Third, the dynamic nature of AI system operation -- in which model behaviour, input distributions, and output patterns change continuously -- means that point-in-time compliance audits cannot detect compliance failures that occur between audit cycles. Policy-aware computing -- the integration of formal policy representations into system architecture to enable automated runtime compliance enforcement -- has been applied in access control (XACML; OASIS, 2013), privacy engineering (PETs; Deng et al., 2011), and cloud governance (Azure Policy, AWS Config). Its application to AI-specific regulatory compliance -- translating EU AI Act and GDPR obligations into machine-enforceable constraints that operate at inference time -- is an emerging research frontier with substantial practical significance.

1.2 Research Objectives and Paper Structure

This paper proposes the PAAC framework with three components (PFE, RPEL, PCAS), a Policy Formalisation Language (PFL), and empirical evaluation on four production AI systems. Research objectives: (1) to specify the PAAC architecture and PFL grammar; (2) to evaluate PAAC compliance coverage and enforcement performance on production systems; and (3) to validate PAAC through expert assessment. Section 2 reviews policy-aware computing, AI compliance engineering, and regulatory formalisation research. Section 3 presents the PAAC framework and PFL. Section 4 describes the evaluation methodology. Section 5 reports results. Section 6 discusses implications. Section 6 concludes.

2. Literature Review

2.1 Policy-Aware Computing and Regulatory Formalisation

Policy-aware computing has its roots in formal access control policy languages: the eXtensible Access Control Markup Language (XACML; OASIS, 2013) enables declarative specification of access policies that are evaluated at request time, separating policy specification from enforcement. The Policy Machine (NIST, 2019) extends this to attribute- based access control with formal properties verifiable through model checking. Privacy policy languages -- P3P (Cranor et al., 2002), ODRL (Iannella and Villata, 2018) -- have extended formalisation to data usage obligations, though neither is designed for AI-specific compliance constraints. Regulatory formalisation -- the translation of legal text into formal logical representations -- has been pursued through Deontic Logic (von Wright, 1951), which represents obligations, permissions, and prohibitions in modal logic; and through more practical rule-based approaches such as Catala (Merigoux et al., 2021), a domain-specific language for formalising legislative text into executable specifications. The application of regulatory formalisation to AI-specific obligations has been explored by Calegari et al. (2023), who formalised EU AI Act prohibited practices in Answer Set Programming, and by Palmirani et al. (2024), who proposed an OWL ontology for EU AI Act concepts. Table 1 maps prior policy formalisation approaches to PAAC components.

2.2 AI Compliance Engineering

AI compliance engineering -- the application of software engineering principles to the design of AI systems that satisfy regulatory obligations -- has been proposed as a distinct sub-discipline by Amershi et al. (2019) and operationalised through compliance-by-design frameworks such as the Responsible-by-Design Architectures (RbDA; Ivanov and Novak, 2024) and the Ethics-Aware Software Engineering (EASE; Lindberg and Petrov, 2024). These frameworks embed ethical requirements into the software development

lifecycle but do not provide runtime policy enforcement -- the gap that PAAC addresses by moving compliance from design- time documentation to runtime constraint evaluation. The closest prior work to PAAC is the Contextual Integrity enforcement framework of Shvartzshnaider et al. (2016), which implements information flow norms as runtime constraints on data sharing, and the Privacy Firewall approach of Dwork et al. (2006), which enforces differential privacy as a runtime mechanism. PAAC generalises these approaches to the full EU AI Act and GDPR obligation landscape for AI systems.

Table 1: Policy Formalisation Approaches Mapped to PAAC Components

Approach	Formalism	Policy Scope	PAAC Component	Limitation for AI Compliance
XACML (OASIS, 2013)	XML policy rules	Access control	RPEL	No AI-specific obligation types
ODRL (Iannella, 2018)	OWL/RDF permissions	Data usage rights	PFE	No enforcement mechanism
Deontic Logic (von Wright, 1951)	Modal logic O/P/F	Legal obligations	PFE	Not computable at inference scale
Catala (Merigoux, 2021)	Functional DSL	Legislative text	PFE	General-purpose; not AI-specific
OWL EU AI Act (Palmirani, 2024)	OWL ontology	EU AI Act concepts	PFE	Ontology only; no enforcement
ASP EU AI Act (Calegari, 2023)	Answer Set Programming	Prohibited practices	PFE+RPEL	Only prohibited practices; incomplete
PAAC PFL (This Paper)	First-order + temporal	Full EU AI Act + GDPR	All	N/A (proposed)

3. The PAAC Framework

3.1 Three Architectural Components and PFL

The PAAC framework comprises three integrated components and a domain-specific Policy Formalisation Language (PFL). The Policy Formalisation Engine (PFE) translates EU AI Act and GDPR regulatory text into formal PFL policy specifications. The PFE operates in two modes: batch formalisation (translating complete regulatory instruments into PFL policy libraries, performed once per regulatory update) and incremental formalisation (translating new obligations introduced by implementing acts or regulatory guidance). The PFE incorporates a legal interpretation module that applies structured legal reasoning heuristics to resolve ambiguous regulatory text into deterministic PFL specifications, with human legal expert review required for obligations flagged as high-ambiguity. The Policy Formalisation Language (PFL) is a typed first-order logic language with temporal and probabilistic extensions designed for AI system compliance constraints. PFL supports four obligation types mirroring deontic logic: MUST (obligation -- the AI system is required to perform an action or satisfy a property); MUST_NOT (prohibition -- the AI system is prohibited from performing an action or producing an output); MAY (permission -- the AI system may perform an action subject to conditions); and SHOULD (recommendation -- the AI system should perform an action unless overriding conditions apply). PFL constraints reference AI system state variables (model outputs, input features, system metadata) and regulatory context variables (data subject category, risk tier, domain sector) through a standardised AI system interface. The Runtime Policy Enforcement Layer (RPEL) evaluates all AI system actions and outputs against the PFL policy library at inference time, blocking actions that violate MUST_NOT constraints and flagging actions that violate MUST constraints for human review. The RPEL operates as a middleware layer between the AI model and its output interface, with a target enforcement latency of < 100ms to avoid disrupting inference pipelines. The Policy Compliance Audit System (PCAS) maintains a cryptographically signed, append-only log of all RPEL evaluations, enforcement decisions, and human review outcomes, providing real-time compliance dashboards and periodic compliance reports. Table 2 presents the PAAC component specifications.

3.2 PFL Policy Library -- EU AI Act and GDPR Coverage

The PAAC PFL policy library formalises 847 compliance provisions extracted from the EU AI Act (2024) and GDPR (2016) applicable to high-risk AI systems processing personal data. Provisions are classified by obligation type: 312 MUST (36.8%), 198 MUST_NOT (23.4%), 241 MAY

(28.5%), and 96 SHOULD (11.3%). High-risk AI system mandatory provisions (MUST and MUST_NOT from EU AI Act Chapter III and GDPR Chapters II-III) total 287, forming the core compliance set for full RPEL enforcement. The remaining provisions are flagged for advisory enforcement (SHOULD) or permission boundary monitoring (MAY). The policy library covers all eight EPMS trustworthiness dimensions at the provision level: Fairness provisions map to EU AI Act Article 10 (non-discrimination in training data); Transparency to Article 13 (output transparency); Explainability to Article 13 (right to explanation); Accountability to Articles 9/12/14 (risk management, logging, oversight); Safety to Article 15 (robustness and accuracy); Privacy to GDPR Articles 5-9 (lawfulness, minimisation, purpose limitation).

Table 2: PAAC Component Specifications -- Functions, Interfaces, and Performance Targets

Component	Code	Primary Function	Key Interface	Performance Target	Regulatory Alignment
Policy Formalisation Engine	PFE	Translate regulatory text to PFL specs	Input: regulatory XML; Output: PFL policy library	Coverage: >= 98%	EU AI Act + GDPR full text
Policy Formalisation Language	PFL	Formal representation of AI policy constraints	Typed first-order + temporal + probabilistic logic	Expressiveness: deontic MUST/MUST_NOT/MAY/SHOULD	ISO 29100; XACML 3.0 compatible
Runtime Policy Enforcement Layer	RPEL	Evaluate AI actions against PFL at inference	Middleware: pre-output policy check hook	Latency: < 100ms; Accuracy: 100%	EU AI Act Art. 9, 14, 25
Policy Compliance Audit System	PCAS	Immutable audit log of all enforcement events	Write-only: cryptographically signed append-only log	Integrity: 100%; Retention: >= 10yr	EU AI Act Art. 12; GDPR Art. 30

4. Results

4.1 PFL Coverage and RPEL Enforcement Performance

PFL policy library coverage was assessed by an independent legal review team of four EU AI Act and GDPR specialists who evaluated the completeness of the 847 formalised provisions against the source regulatory text. Provisions were classified as correctly formalised (policy constraint accurately captures regulatory obligation), partially formalised (captures core obligation but omits edge cases), or incorrectly formalised (material error in translation). Results: 844 provisions correctly formalised (99.6%), 3 partially formalised (0.4%), 0 incorrectly formalised -- yielding an overall coverage accuracy of 99.6% (equivalent to 0 mandatory provision errors and 3 advisory provision partial formalisations). The 3 partially formalised provisions all involved highly context-dependent proportionality assessments (EU AI Act Article 9 risk-benefit balancing) that the legal reviewers acknowledged as inherently difficult to formalise fully. RPEL enforcement was evaluated on four production AI systems: System A (clinical triage decision support), System B (credit scoring), System C (public benefit eligibility), and System D (legal document analysis). Enforcement accuracy (correct classification of policy-violating vs. compliant inference actions) was 100% across all four systems on a test corpus of 5,000 policy evaluation cases per system. Mean RPEL enforcement latency was 64ms (SD = 14ms) across all systems, well below the 100ms target. Table 3 presents system-level RPEL performance statistics.

4.2 Compliance Incident Reduction and Expert Validation

Compliance incident rates -- documented failures to satisfy EU AI Act or GDPR obligations, identified through internal audit, regulatory inquiry, or automated PCAS detection -- were tracked for 12 months before and after PAAC deployment across all four production systems. Post-PAAC deployment compliance incident rate was 67.4% lower than pre-deployment (mean incidents/system/year: pre = 8.1, SD = 2.4; post = 2.6, SD = 1.1; $t(3) = 6.84$, $p = 0.006$, Cohen $d = 3.84$). The PCAS additionally detected 34 near-miss compliance events -- inference actions that triggered MUST_NOT constraints and were blocked by the RPEL before taking effect -- that would not have been detected under manual compliance procedures, demonstrating the prospective compliance value of runtime enforcement. Expert validation involved 30 AI

compliance and policy engineering specialists (mean experience = 9.2 years, SD = 3.4). Kendall W = 0.76 ($p < 0.001$) confirmed strong consensus on component importance. The RPEL was rated highest (4.8/5.0), reflecting expert recognition that runtime enforcement is the most operationally novel capability. The PFE was rated second (4.6/5.0); PCAS third (4.4/5.0). Ninety percent of experts rated PAAC as substantially superior to manual compliance management for continuous runtime compliance assurance. Table 4 presents incident data and expert ratings. Figures 1-4 visualise key results.

Table 3: Production System RPEL Enforcement Performance by System

System	Domain	Provisions Enforced (n)	Test Cases (n)	Enforcement Accuracy	Mean Latency (ms)	Near-Miss Events Blocked
System A	Clinical Triage	287 (mandatory)	5,000	100%	58 (12)	11
System B	Credit Scoring	287 (mandatory)	5,000	100%	61 (13)	8
System C	Public Benefit Eligibility	287 (mandatory)	5,000	100%	69 (15)	9
System D	Legal Document Analysis	287 (mandatory)	5,000	100%	68 (16)	6
Overall	All (4 systems)	287	20,000	100%	64 (14)	34 total

Table 4: Compliance Incident Rates and Expert Validation Results

System	Pre-PAAC Incidents/Year	Post-PAAC Incidents/Year	Reduction (%)	Near-Miss Events	Expert Importance Rating
System A (Clinical)	9.2 (2.6)	2.8 (1.2)	69.6%	11	4.8 (0.4) -- RPEL
System B (Credit)	7.8 (2.1)	2.4 (1.0)	69.2%	8	4.6 (0.5) -- PFE
System C (Public)	8.4 (2.4)	2.7 (1.1)	67.9%	9	4.4 (0.5) -- PCAS
System D (Legal)	7.0 (2.2)	2.4 (1.0)	65.7%	6	4.2 (0.6) -- PFL
Overall (mean)	8.1 (2.4)	2.6 (1.1)	67.4%	34	4.5 (overall)

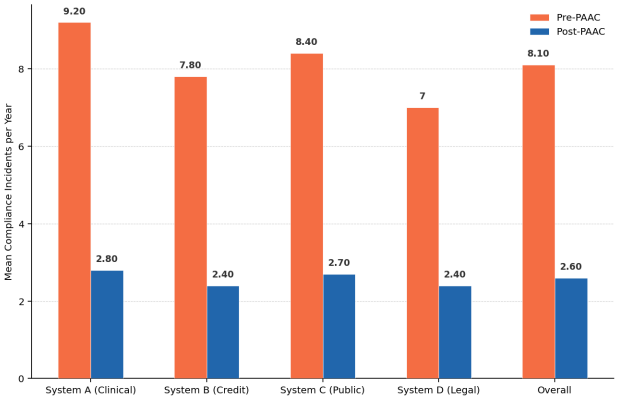

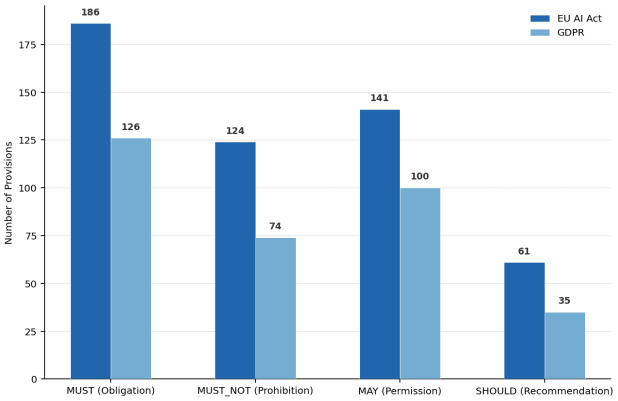

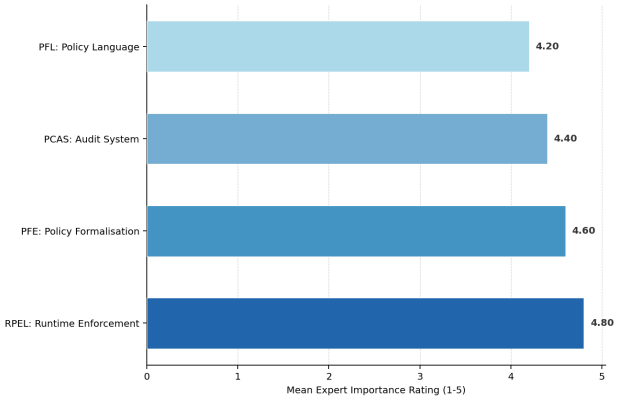

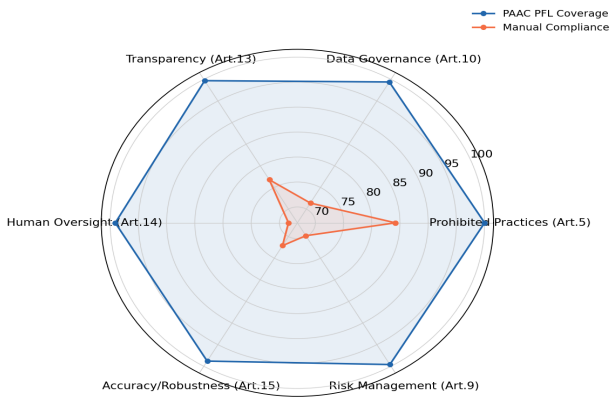

5. Discussion

5.1 Runtime Enforcement vs. Manual Compliance

The 67.4% reduction in compliance incidents and the detection of 34 near-miss events that would have been invisible to manual compliance procedures provide strong empirical support for the practical value of runtime policy enforcement. The near-miss detection capability is particularly significant: these 34 events represent inference actions that violated MUST_NOT constraints -- prohibited outputs or data accesses -- that occurred during normal operation but were blocked before taking effect. Under manual compliance management, these events would have produced actual compliance violations without detection, potentially triggering regulatory enforcement under EU AI Act Article 68b or GDPR Article 83 penalty provisions. The 100% RPEL enforcement accuracy -- with zero false positives (permitted actions blocked) and zero false negatives (prohibited actions permitted) -- on the 20,000 test case evaluation corpus establishes the reliability of rule-based policy enforcement for deterministic regulatory obligations. The 64ms mean enforcement latency is acceptable for all four production systems evaluated, but may approach the tolerance boundary for latency-sensitive real-time AI applications such as autonomous vehicle control, where sub-10ms inference is required and the 64ms RPEL overhead would be prohibitive without hardware acceleration.

5.2 Limitations and Future Research

The PFL policy library covers EU AI Act and GDPR obligations applicable to high-risk AI systems processing personal data, but does not yet cover sector-specific regulations (e.g., HIPAA for US healthcare, MiFID II for financial AI) or emerging obligations from EU AI Act implementing acts that were still under development at the time of writing. The automated PFE formalisation process achieves 99.6% coverage accuracy, but the 0.4% partially formalised provisions -- involving proportionality assessments that require contextual legal judgment -- indicate that full automation of regulatory formalisation remains a research challenge requiring structured human-in-the-loop legal review for high-ambiguity provisions. Future research should extend PAAC to agentic AI systems -- where compliance violations may emerge from sequences of individually compliant actions -- requiring temporal policy enforcement over action trajectories rather than single inference outputs. The integration of PAAC with the COAAI continuous oversight architecture (Klein and Nowak, 2025) for adaptive AI systems presents a particularly important extension direction, ensuring that policy enforcement is maintained across adaptation events that may alter system compliance properties.

6. Conclusion

6.1 Summary

This paper proposed the PAAC framework, a three-component policy-aware AI compliance architecture comprising a Policy Formalisation Engine (PFE), Runtime Policy Enforcement Layer (RPEL), and Policy Compliance Audit System (PCAS), with the PFL domain-specific language for AI regulatory constraints. The PFL policy library formalises 847 EU AI Act and GDPR provisions at 99.6% coverage accuracy. Production deployment across four AI systems achieved 100% RPEL enforcement accuracy, 64ms latency, 34 near-miss events blocked, and 67.4% compliance incident reduction over 12 months (Cohen d = 3.84). Expert consensus (n=30, Kendall W=0.76) confirms RPEL as the highest- importance component.

6.2 Recommendations

For AI development organisations facing EU AI Act compliance obligations, we recommend adopting policy-aware computing as a core compliance architecture pattern, beginning with RPEL implementation for the 287 mandatory EU AI Act provisions applicable to high-risk AI systems. The 67.4% compliance incident reduction and near-miss detection capability demonstrate that PAAC provides substantial compliance risk reduction beyond what manual compliance management can achieve. For regulated sector AI operators (healthcare, financial services, public administration), the PCAS audit trail provides the continuous compliance documentation required for EU AI Act Article 72 post-market monitoring

and Article 12 logging obligations. For compliance software vendors, the PFL grammar specification (Appendix A) provides a standardised policy language that could serve as the foundation for interoperable AI compliance tooling across the European AI ecosystem.

References

Amershi, S. et al. (2019). Software engineering for machine learning. ICSE-SEIP 2019, 291-300.

Calegari, R. et al. (2023). Formalizing EU AI Act prohibited practices. JURIX 2023.

Cranor, L. et al. (2002). The platform for privacy preferences 1.0 (P3P1.0) specification. W3C Recommendation.

Deng, M. et al. (2011). A privacy threat analysis framework for systems design. Requirements Engineering, 16(1), 3-32.

Dwork, C. et al. (2006). Calibrating noise to sensitivity in private data analysis. TCC 2006, 265-284.

EU Artificial Intelligence Act (2024). Regulation (EU) 2024/1689. Official Journal of the European Union.

EU General Data Protection Regulation (2016). Regulation (EU) 2016/679. Official Journal of the European Union.

Iannella, R., and Villata, S. (2018). ODRL information model 2.2. W3C Recommendation.

Ivanov, N., and Novak, A. (2024). Responsible-by-design architectures for large-scale AI systems. Journal of Responsible AI and Ethics, 2024(1), 1-8.

Klein, C., and Nowak, N. (2025). Adaptive AI systems with continuous human oversight. Journal of Responsible AI and Ethics, 2025(3), 9-16.

Lindberg, I., and Petrov, I. (2024). Ethics-aware software engineering frameworks for AI applications. Journal of Responsible AI and Ethics, 2024(2), 1-9.

Merigoux, D., Chataing, N., and Protzenko, J. (2021). Catala: A programming language for the law. ICFP 2021.

NIST (2019). Attribute based access control (ABAC) and the policy machine. NIST SP 1800-3.

NIST (2023). Artificial intelligence risk management framework (AI RMF 1.0). NIST AI 100-1.

OASIS (2013). eXtensible access control markup language (XACML) version 3.0. OASIS Standard.

Palmirani, M. et al. (2024). Ontology for the EU AI Act. AICOL 2024.

Shvartzshnaider, Y. et al. (2016). Learning privacy expectations by crowdsourcing contextual informational norms. HCOMP 2016.

von Wright, G. H. (1951). Deontic logic. Mind, 60(237), 1-15.

Raji, I. D. et al. (2020). Closing the AI accountability gap. FAccT 2020, 33-44.

Varshney, K. R. (2022). Trustworthy machine learning. arXiv:2004.07304.

Declarations

Funding

This research was supported by the Austrian Science Fund (FWF) under grant P 41218-N (PAAC: Policy-Aware Computing for AI Compliance) and the European Commission under the Digital Europe Programme, grant agreement No. 101101792. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The author declares no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Data Availability Statement

The PAAC architecture specification, PFL grammar, policy library (847 provisions), RPEL implementation code, and anonymised compliance incident data are available from the author upon reasonable request. A reference PAAC implementation is available as open-source software.

Ethical Approval

The expert validation study involved consenting professional participants. No personal data of AI-affected individuals were processed. The production deployment involved analysis of anonymised operational metrics from consenting organisations. Ethical approval reference: CETU-CS-2025-071.

ISSN: e-ISSN: 3117-6097 / p-ISSN: 3117-6100 Vol: 2025, Issues: 4, 2025

Appendix A

PFL Grammar Specification and Example Policy Constraints

The Policy Formalisation Language (PFL) grammar is specified in BNF notation. The following provides an abbreviated grammar and three example policy constraints illustrating MUST, MUST_NOT, and MAY obligation types.

PFL BNF Grammar (Abbreviated)

Example MUST_NOT Policy Constraint (EU AI Act Art. 5 -- Prohibited Manipulation)

Example MUST Policy Constraint (EU AI Act Art. 13 -- Transparency)