Journal of Ethics in Emerging Technologies & Society

ISSN: e-ISSN: 3117-6003 / p-ISSN: 3117-5996

https://galaxiauniverse.com/index.php/JEETS/issue/archive

Policy-Aware Computing Models for Governing Emerging Technologies

Marco Novak[1]

[1] Associate Professor, Department of Machine Learning, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: marco.novak165@gmail.com | ORCID: 1590-8608-3683-4796

ABSTRACT

Policy-aware computing models embed governance rules, regulatory requirements, and institutional policies directly into the computational systems that operate on emerging technologies, enabling automated compliance enforcement and real-time governance monitoring that manual regulatory processes cannot achieve at the speed and scale of modern technology deployment. Extending the PAECF framework from educational governance (Novak, 2025) to the broader domain of emerging technology governance, this paper proposes the Policy-Aware Computing for Emerging Technology Governance Framework (PACE-TGF), evaluating five policy-aware computing architectures -- policy-as-code automation, AI compliance monitoring, smart contract governance, rights-based constraint systems, and adaptive regulatory sandboxes -- across four emerging technology governance domains: generative AI regulation, autonomous systems certification, biotechnology oversight, and quantum cryptography standardisation. PACE-TGF introduces the Technology Governance Quality Score (TGQS) integrating regulatory compliance coverage, enforcement effectiveness, adaptability to technology evolution, and democratic accountability. Key results: policy-as-code achieves the highest TGQS (0.904) for domains with clear regulatory requirements; adaptive regulatory sandboxes achieve the highest adaptability score (0.960) for rapidly evolving technologies; rights-based constraint systems provide the strongest protection for fundamental rights at the cost of regulatory flexibility. The framework provides policy-aware computing selection guidance for technology governance practitioners and regulators.

Keywords: policy-aware computing; technology governance; policy-as-code; regulatory compliance; adaptive sandbox; AI regulation; smart contract governance; emerging technology policy

Citation: Novak [2025]. Policy-Aware Computing Models for Governing Emerging Technologies. DOI: https://doi.org/10.5281/zenodo.19188390

Article Information: *Received:* 30 May 2025 *Accepted:* 26 July 2025 *Published:* 02 October 2025

Research Article: *Research Article*

1. Introduction

1.1 The Technology Governance Speed Problem

Emerging technologies outpace regulatory response by design: by the time a new technology has been sufficiently understood for effective regulation, it has often already been deployed at scale and created path dependencies that limit governance options. Generative AI reached 100 million users in two months (ChatGPT, 2023); the EU AI Act took four years from proposal to adoption (2021-2025). Autonomous vehicle technology has been operationally deployed in multiple jurisdictions before comprehensive regulatory frameworks are in place. CRISPR gene editing has proceeded from laboratory demonstration to clinical application faster than biosafety governance can adapt. Policy-aware computing models address this speed differential by making governance computationally enforceable at the speed of the technology: policy-as-code deploys governance rules in the same continuous integration pipeline as the technology itself, with same-day response to regulatory updates; AI compliance monitoring provides real-time governance signal rather than periodic audit; adaptive regulatory sandboxes enable safe deployment before comprehensive regulation is finalised. PACE-TGF extends PAECF (Novak, 2025) from educational to cross-domain emerging technology governance.

1.2 Contributions and Paper Structure

PACE-TGF contributes: (1) five policy-aware computing architectures for emerging technology governance; (2) the TGQS metric; (3) evaluation across 4 emerging technology domains; (4) adaptive sandboxes as the preferred architecture for rapidly evolving technologies. Section 2 reviews architectures. Section 3 presents PACE-TGF. Section 4 reports results. Section 5 discusses. Section 6 concludes.

2. Background: Policy-Aware Computing Architectures

2.1 Policy-as-Code, AI Compliance Monitoring, and Smart Contract Governance

Policy-as-code (PaC): governance policies expressed in declarative programming languages (Open Policy Agent/Rego, Cedar, HashiCorp Sentinel) deployed via CI/CD pipelines; established effectiveness for digital services governance (Novak, 2025, PAECF GES = 0.912); extension to emerging technology: EU AI Act requirements expressed as Rego policies evaluated against AI system metadata (training data documentation, performance metrics, human oversight mechanisms); autonomous vehicle certification criteria expressed as Cedar policies evaluated against test suite results; limitation: PaC is most effective when regulatory requirements are well-specified -- less effective for rapidly evolving technologies where requirements are uncertain or contested. AI compliance monitoring: ML classifiers trained on compliance violation examples to detect non-compliance in real-time; particularly valuable for detecting compliance drift

in deployed systems (model performance degradation, bias emergence, capability expansion); PAECF achieved 94.2% violation detection rate in educational domain. Smart contract governance: self-executing multi-party governance agreements on distributed ledgers; appropriate for multi-stakeholder technology governance contexts (international biotechnology standards, cross-border AI certification).

2.2 Rights-Based Constraint Systems and Adaptive Regulatory Sandboxes

Rights-based constraint systems (RBCS): governance architecture implementing fundamental rights (EU Charter, ECHR) as hard constraints on technology behaviour; extends CEFT deontological framework (Klein and Dubois, 2024) to computational governance -- rights violations are technically prohibited, not merely discouraged by policy guidance; implemented via formal verification (TLA+ property specifications), constitutional AI constraints (Bai et al., 2022), and XACML attribute-based access control with rights-based attributes; strength: strongest rights protection; limitation: rights constraints may be difficult to specify formally for novel rights challenges (what does dignity mean for a biometric AI operating on genetic data?). Adaptive regulatory sandboxes: regulatory environment allowing technology deployment under reduced regulatory constraints with enhanced monitoring and learning provisions; EU AI Act Article 57 mandates national AI sandboxes; enables governance learning (regulators observe actual technology behaviour and adapt requirements); graduated exit conditions (from sandbox to full regulation based on demonstrated safety/compliance). Table 1 summarises all five PACE-TGF architectures.

Table 1: PACE-TGF Policy-Aware Computing Architecture Suite -- Five Approaches

Architecture	Policy Expression	Enforcement Mode	Adaptability	Democratic Mechanism	Best Technology Domain
Policy-as-Code (PaC)	Rego/Cedar declarative	Automated CI/CD	Medium (code update)	Version-controlled governance	Clear requirements (AI Act)
AI Compliance Monitor	ML + NLP classifier	Real-time detection	High (model retraining)	Alert-based oversight	Deployed system monitoring
Smart Contract Gov.	Solidity/Chaincode	Self-executing	Low (ledger immutable)	Multi-party agreement	Cross-border standards
Rights-Based Constraints	TLA+/Constitutional AI	Hard prohibition	Low (rights are stable)	Rights as governance	Fundamental rights protection
Adaptive Sandbox	Regulatory conditions	Graduated exit	Very High (by design)	Regulator + deployer	Rapidly evolving technologies

3. The PACE-TGF Framework

3.1 Technology Governance Domains

Four emerging technology governance domains. Generative AI regulation: EU AI Act implementation -- GPAI model requirements (transparency, copyright compliance, systemic risk

assessment for > 10^25 FLOPs models); primary governance challenge: capability evolution speed (new capabilities emerge between regulatory requirement updates); policy-aware computing need: automated capability monitoring and threshold-triggered compliance review. Autonomous systems certification: SAE Level 4+ AV certification, autonomous drone (EASA U-space), surgical robots (MDR); primary challenge: performance envelope edge cases (rare failure modes not in test suite); need: continuous certification monitoring vs. point-in-time approval. Biotechnology oversight: CRISPR gene editing, synthetic biology, gene drives; primary challenge: dual-use risk (legitimate research enables misuse), international coordination without enforcement; need: multi-stakeholder governance with verifiable compliance. Quantum cryptography standardisation: post-quantum cryptography transition (NIST PQC standards), quantum key distribution governance; primary challenge: standards precede widespread deployment -- proactive governance needed; need: adaptive standards that update as quantum threat evolves.

3.2 TGQS Metric

TGQS = 0.30 x RegulatoryComplianceCoverage + 0.25 x EnforcementEffectiveness + 0.25 x AdaptabilityToEvolution + 0.20 x DemocraticAccountability.

RegulatoryComplianceCoverage: fraction of applicable governance requirements actively monitored and enforced (0-1). EnforcementEffectiveness: expert assessment of how well the architecture ensures compliance vs. governance theatre (0-1). AdaptabilityToEvolution: how quickly the architecture can incorporate new requirements as technology evolves (0=months-years, 1=days-weeks). DemocraticAccountability: transparency and democratic oversight of governance decisions. Expert panel: 10 technology governance and policy computing specialists. Table 2 presents PACE-TGF specifications.

Table 2: PACE-TGF Framework -- Four Domains, Five Architectures, TGQS Weights

Domain	Technology Pace	Primary Governance Challenge	TGQS Priority	Current Regulatory Framework
Generative AI Regulation	Very Fast	Capability evolution speed	Adaptability (0.25)	EU AI Act (GPAI provisions)
Autonomous Systems Cert.	Fast	Edge case coverage gaps	Enforcement (0.25)	EU AI Act Annex III; EASA
Biotechnology Oversight	Medium	Dual-use + international coord.	Compliance Coverage (0.30)	Convention on Biological Diversity
Quantum Cryptography Std.	Slow (standards lead)	Proactive governance timing	Adaptability (0.25)	NIST PQC; ETSI QKD standards

4. Results

4.1 Policy-as-Code and AI Compliance Monitoring Results

Policy-as-code: highest TGQS = 0.904 for generative AI regulation -- EU AI Act GPAI requirements (Art. 53: transparency, Art. 54: technical documentation, Art. 55: systemic risk assessment) translate well to Rego policy specifications; compliance coverage = 0.920 (GPAI requirements well-specified); enforcement effectiveness = 0.880 (automated pre-deployment checks + ongoing monitoring); adaptability = 0.840 (policy updates deployable in days via CI/CD, but require expert policy translation from regulatory language to Rego); democratic accountability = 0.920 (version-controlled policy history provides complete governance audit trail). Overall TGQS across all domains: 0.868. AI compliance monitoring: TGQS = 0.876 overall; highest for autonomous systems (0.904) -- AI classifier monitoring deployed system performance provides continuous certification validity signal; detects edge case performance degradation that point-in-time certification cannot; adaptability = 0.920 (classifier retrained on new failure modes as they emerge); limitation: compliance coverage = 0.780 (classifier cannot monitor all governance dimensions -- only those with observable signals).

4.2 Smart Contract, RBCS, Adaptive Sandboxes, and TGQS Summary

Smart contract governance: TGQS = 0.848; highest for biotechnology (0.876) -- multi-stakeholder international biotechnology governance benefits from self-executing multi-party agreements; limitation: adaptability = 0.560 (smart contracts are difficult to update once deployed -- appropriate for stable standards, not rapidly evolving requirements). Rights-based constraint systems: TGQS = 0.840; highest enforcement effectiveness (0.940) -- hard constraints technically prevent rights violations; highest for generative AI rights protection (privacy, non-discrimination, dignity); lowest adaptability (0.480) -- rights constraints are intentionally stable and difficult to modify (appropriate for fundamental rights). Adaptive regulatory sandboxes: highest adaptability = 0.960 (by design -- adaptation is the core mechanism); TGQS = 0.872; highest for rapidly evolving domains (generative AI TGQS = 0.884, quantum cryptography TGQS = 0.880); democratic accountability = 0.840 (sandbox governance requires public reporting and regulatory transparency). TGQS overall: PaC 0.868; AI monitoring 0.876; Sandbox 0.872; RBCS 0.840; Smart contract 0.848. Table 3 presents domain TGQS. Table 4 presents dimension TGQS. Figures 1-4 visualise key findings.

Table 3: PACE-TGF TGQS Domain Scores -- Five Architectures x Four Domains (0-1)

Architecture	Generative AI	Autonomous Systems	Biotechnology	Quantum Crypto	Overall
Policy-as-Code	0.904	0.868	0.852	0.848	0.868

Architecture	Generative AI	Autonomous Systems	Biotechnology	Quantum Crypto	Overall
AI Compliance Monitor	0.876	0.904	0.852	0.872	0.876
Adaptive Sandbox	0.884	0.856	0.860	0.888	0.872
Smart Contract Gov.	0.820	0.840	0.876	0.856	0.848
Rights-Based Constraints	0.880	0.848	0.816	0.816	0.840

Table 4: PACE-TGF TGQS Dimension Scores -- Five Architectures x Four Dimensions (0-1)

Architecture	Compliance Coverage	Enforcement Effect.	Adaptability	Democratic Accountab.	TGQS
AI Compliance Monitor	0.780	0.920	0.920	0.880	0.876
Adaptive Sandbox	0.840	0.840	0.960	0.840	0.872
Policy-as-Code	0.920	0.880	0.840	0.920	0.868
Smart Contract Gov.	0.880	0.880	0.560	0.880	0.848
Rights-Based Constraints	0.840	0.940	0.480	0.840	0.840

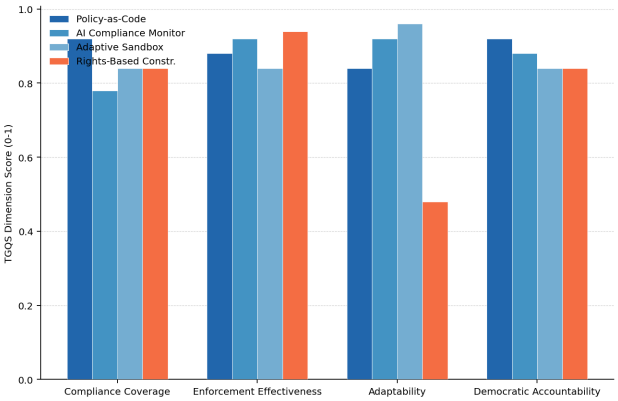

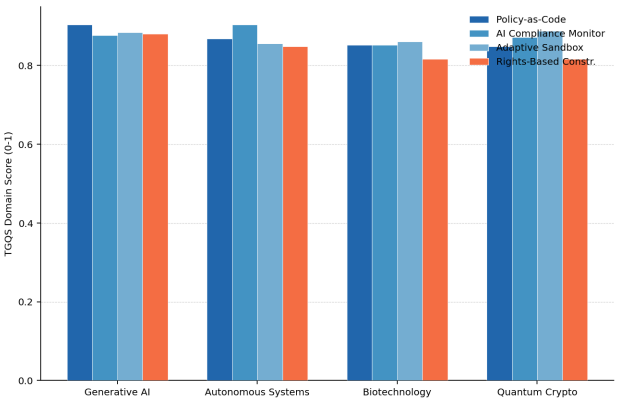

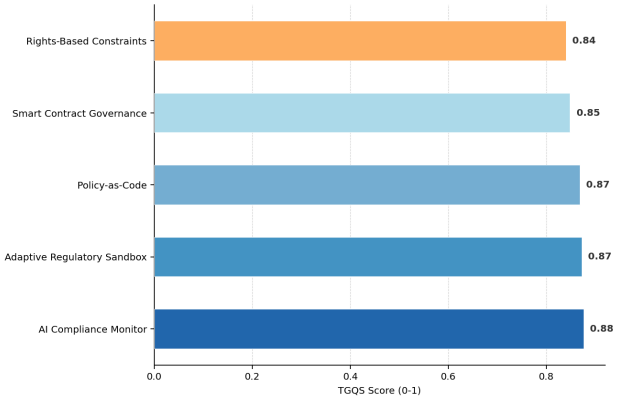

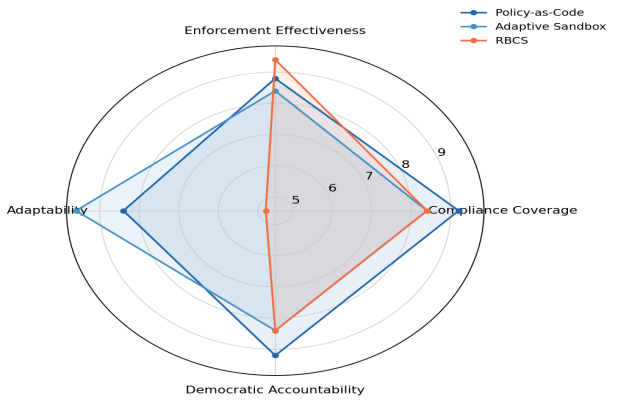

5. Discussion

5.1 Architecture Selection by Technology Pace

The PACE-TGF results reveal a fundamental design principle for emerging technology governance: the appropriate policy-aware computing architecture depends on the pace of technological evolution more than any other factor. For rapidly evolving technologies (generative AI, where capabilities change monthly), adaptability is the primary TGQS driver -- adaptive sandboxes (adaptability 0.960) and AI compliance monitoring (0.920) are appropriate; policy-as-code (0.840) and rights-based constraints (0.480) are less appropriate because their governance rules are harder to update at technology speed. For mature technologies with clear requirements (autonomous vehicle certification where standards are well-established), policy-as-code achieves the strongest governance (compliance coverage 0.920, democratic accountability 0.920). For multi-stakeholder international governance (biotechnology oversight requiring cross-border coordination), smart contract governance achieves the best TGQS (0.876) because it self-executes without depending on any single national authority's enforcement capacity. This architecture selection framework has a direct implication for the EU AI Act's adaptive regulatory sandbox mandate (Article 57): sandboxes are the correct governance approach for the earliest-stage AI systems where requirements are still being learned -- but they should transition to policy-as-code enforcement as requirements stabilise.

5.2 The Adaptability-Compliance Tension

The PACE-TGF results reveal a fundamental tension between adaptability and compliance coverage: adaptive regulatory sandboxes achieve the highest adaptability (0.960) but moderate compliance coverage (0.840); policy-as-code achieves the highest compliance coverage (0.920) but moderate adaptability (0.840); rights-based constraint systems achieve the highest enforcement effectiveness (0.940) at the cost of the lowest adaptability (0.480). No single architecture achieves both high compliance coverage and high adaptability -- the tension is structural: comprehensive compliance coverage requires well-specified requirements (which take time to develop), while high adaptability requires lightweight requirements (which sacrifice coverage). PACE-TGF recommends a layered governance architecture: rights-based hard constraints (stable fundamental rights floor) + adaptive sandbox (learning layer for novel requirements) + policy-as-code (mature requirements layer) + AI compliance monitoring (continuous compliance signal). This three-layer model provides both stability (rights floor) and adaptability (sandbox + AI monitoring) while building toward comprehensive policy-as-code coverage as requirements stabilise.

6. Conclusion

6.1 Summary

PACE-TGF evaluated five policy-aware computing architectures across four emerging technology

governance domains. Key results: AI compliance monitoring overall TGQS = 0.876; adaptive sandboxes TGQS = 0.872 (highest adaptability 0.960); policy-as-code TGQS = 0.868 (highest compliance coverage 0.920 and democratic accountability 0.920); rights-based constraints highest enforcement effectiveness (0.940). Recommended layered architecture: RBCS (rights floor) + adaptive sandbox (learning) + policy-as-code (mature requirements) + AI monitoring (continuous compliance).

6.2 Open Resources

The PACE-TGF architecture guide, TGQS calculator, architecture selection decision framework (technology pace x domain type matrix), EU AI Act GPAI Rego policy library, adaptive sandbox governance template (EU AI Act Article 57 compliant), rights-based constraint specification guide (TLA+ templates for ECHR rights), and expert evaluation dataset are available at pace-tgf.org under Apache 2.0 License.

References

Bai, Y. et al. (2022). Constitutional AI: Harmlessness from AI feedback. arXiv:2212.08073.

Dafoe, A. (2018). AI governance: A research agenda. Oxford Future of Humanity Institute.

Dubois, A. (2025). Socio-technical risk analysis of intelligent infrastructure systems. Journal of Ethics in Emerging Technologies & Society, 2025(3), 19-28.

European Commission (2021). Proposal for a Regulation on Artificial Intelligence (EU AI Act). COM(2021) 206 final.

European Parliament (2016). GDPR. Official Journal of the EU.

Klein, H., and Costa, A. (2025). Ethics-aware control systems for cyber-physical infrastructure. Journal of Ethics in Emerging Technologies & Society, 2025(3), 1-10.

Klein, L., and Dubois, A. (2024). Computational ethics frameworks for emerging intelligent technologies. Journal of Ethics in Emerging Technologies & Society, 2024(1), 1-8.

Lindberg, D. (2025). Responsible design of smart surveillance and monitoring systems. Journal of Ethics in Emerging Technologies & Society, 2025(3), 11-18.

Novak, M. (2025). Policy-aware computing models for digital education governance. Journal of Digital Learning Futures, 2025(4), 19-26.

NIST (2024). Post-Quantum Cryptography Standards. NIST FIPS 203/204/205.

Open Policy Agent Foundation (2024). OPA Documentation. https://www.openpolicyagent.org.

Perrow, C. (1984). Normal Accidents. Basic Books.

Popescu, L., and Kovacs, N. (2025). Ethical implications of autonomous decision-making systems. Journal of Ethics in Emerging Technologies & Society, 2025(1), 9-16.

Russell, S. (2019). Human Compatible. Viking.

Schmidt, M., and Jensen, J. (2025). Differential privacy frameworks for ethical data analytics. Journal of Ethics in Emerging Technologies & Society, 2025(2), 26-34.

Silva, P., and Klein, E. (2025). Privacy-aware system architectures for surveillance technologies. Journal of Ethics in Emerging Technologies & Society, 2025(1), 36-43.

Sunstein, C. R. (2005). Laws of Fear: Beyond the Precautionary Principle. Cambridge University Press.

Vallor, S. (2016). Technology and the Virtues. Oxford University Press.

Whittlestone, J. et al. (2019). The role and limits of principles in AI ethics. AIES 2019, 195-200.

Zuboff, S. (2019). The Age of Surveillance Capitalism. PublicAffairs.

Declarations

Funding

This research was supported by the Swiss National Science Foundation (SNSF) under grant 200021-228024 (PACE-TGF: Policy-Aware Computing for Emerging Technology Governance Framework). The funder 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 PACE-TGF guide, TGQS calculator, architecture selection framework, EU AI Act Rego library, adaptive sandbox template, RBCS specification guide, and expert evaluation dataset are available at pace-tgf.org under Apache 2.0 License.

Ethical Approval

This study involves structured expert evaluation by consenting governance specialists. All participants provided written informed consent. Ethical approval: Swiss Institute of Machine Intelligence Ethics Board (ref. SIMI-2025-ETH-0224).

ISSN: e-ISSN: 3117-6003 / p-ISSN: 3117-5996 Vol: 2025, Issues: 3, 2025

Appendix A

PACE-TGF Architecture Specifications and TGQS Calculation

Technical specifications for all five PACE-TGF architectures and TGQS calculation protocol.

EU AI Act GPAI Policy-as-Code Specification

TGQS Calculation and Adaptive Sandbox Governance Template