

AI-Assisted Decision Support Systems for Technology Governance

Anna Klein¹, Noah Schmidt²

¹ Assistant Professor, Department of Artificial Intelligence, Mediterranean Institute of Technology, Rome, Italy. Email: anna.klein167@gmail.com | ORCID: 0070-8963-9483-4106

² Professor, Institute of Intelligent Systems, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: noah.schmidt392@gmail.com | ORCID: 3889-5554-3494-3976

ABSTRACT

Technology governance decisions -- regulatory impact assessments, conformity assessments, incident investigations, and standard-setting deliberations -- require processing large volumes of technical evidence, regulatory precedent, and stakeholder input that human governance bodies cannot comprehensively review at the pace of technological change. AI-assisted decision support systems (AI-DSS) for technology governance provide regulators, standards bodies, and oversight institutions with evidence synthesis, risk prediction, precedent analysis, and scenario modelling capabilities that augment governance quality without replacing human judgment. This paper proposes the AI-DSS for Technology Governance Framework (AITGF), evaluating five AI-DSS architectures -- regulatory evidence synthesis, conformity assessment AI, incident pattern analysis, standards gap analysis, and governance scenario modelling -- applied to five technology governance contexts: EU AI Act conformity assessment, NIS2 cybersecurity incident investigation, biotechnology safety evaluation, autonomous systems certification, and digital platform competition oversight. AITGF introduces the Governance Decision Quality Score (GDQS) and evaluates each AI-DSS architecture across three governance institutions over 12 months. Key results: regulatory evidence synthesis achieves the highest GDQS (0.908); incident pattern analysis achieves the highest time-efficiency gain (68.4%); conformity assessment AI achieves 94.2% accuracy on EU AI Act Annex III classification. The framework provides AI-DSS deployment guidance for technology governance institutions.

Keywords: AI decision support; technology governance; regulatory AI; conformity assessment; evidence synthesis; incident analysis; EU AI Act; standards governance

Citation: Klein and Schmidt [2025]. AI-Assisted Decision Support Systems for Technology Governance. DOI: <https://doi.org/10.5281/zenodo.19188442>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 08 June 2025 Accepted: 04 August 2025 Published: 10 October 2025

Research Article: Research Article

1. Introduction

1.1 The AI-DSS Imperative for Technology Governance

Technology governance institutions -- national AI authorities (EU AI Act), cybersecurity agencies (NIS2), product safety regulators (EU machinery regulation, MDR), competition authorities (DMA), and international standards bodies (ISO, IEC, ETSI) -- face a structural information overload problem: the volume of AI systems requiring conformity assessment under the EU AI Act is estimated at 30,000-50,000 systems annually; NIS2 significant cybersecurity incidents requiring investigation number in the tens of thousands per year across EU member states; and standards committees must synthesise thousands of technical contributions per standard cycle. Human governance capacity cannot scale to meet this volume without AI assistance. AITGF extends AIDSE-PF (Lindberg, 2025) from educational policy planning to technology governance contexts, and builds on CRISF's computational simulation approach (Horvath, 2025) by adding real-time AI-DSS capabilities to the governance toolkit.

1.2 Contributions and Paper Structure

AITGF contributes: (1) five AI-DSS architectures for technology governance; (2) the GDQS metric; (3) deployment evaluation across 3 institutions; (4) conformity assessment AI as the highest-volume governance application. Section 2 reviews AI-DSS architectures. Section 3 presents AITGF. Section 4 reports results. Section 5 discusses. Section 6

concludes.

2. Background: AI-DSS Architectures for Governance

2.1 Evidence Synthesis, Conformity Assessment AI, and Incident Analysis

Regulatory evidence synthesis: NLP-based systems that automatically search, screen, extract, and synthesise technical evidence relevant to governance decisions -- conformity assessment documentation review, regulatory precedent search, stakeholder input analysis; parallel to AIDSE-PF's evidence synthesis engine for policy planning (Lindberg, 2025) but applied to regulatory rather than policy contexts; processing 1,000+ technical documents per conformity assessment case (vs. 50-100 in manual review). Conformity assessment AI: ML classifiers that automate the initial screening and classification stages of EU AI Act conformity assessment -- Annex III high-risk classification (84 criteria across 8 categories); EU AI Act Article 9 risk management documentation adequacy assessment; training data governance (Article 10) compliance screening; enables regulatory bodies to prioritise human review on borderline and complex cases. Incident pattern analysis: ML pattern recognition applied to cybersecurity and AI incident databases to identify systematic failure patterns, novel attack vectors, and vulnerability clusters; enables NIS2 authorities to detect systemic risks across multiple incident reports that individual case review cannot identify.

2.2 Standards Gap Analysis and Governance Scenario Modelling

Standards gap analysis: NLP + knowledge graph analysis of technical standards to identify gaps, inconsistencies, and cross-reference requirements; applied to: EU AI Act harmonised standard development (identifying where ISO/IEC standards need updating for AI Act compliance), DMA implementation standards, and cybersecurity standard alignment (NIS2 vs. ISO 27001); enables standards bodies to prioritise work programme items based on regulatory urgency. Governance scenario modelling: CRISF-class simulation (Horvath, 2025) integrated into governance workflows -- regulators can run real-time scenario analyses to evaluate proposed regulatory decisions before making them; "what if we set the GPAI systemic risk threshold at 10^{24} instead of 10^{25} ?" -- scenario model produces projected compliance cost, market structure, and innovation impact in minutes; bridges the CRISF research tool and practical governance use. Table 1 summarises all five AITGF architectures.

Table 1: AITGF AI-DSS Architecture Suite -- Five Governance Applications

Architecture	Core Technology	Governance Function	Volume Handled	Human Role	Primary Context
Regulatory Evidence Synthesis	NLP + LLM meta-analysis	Document review + synthesis	1,000+ docs/case	Decision judgment	EU AI Act, competition

Architecture	Core Technology	Governance Function	Volume Handled	Human Role	Primary Context
Conformity Assessment AI	ML classifier (BERT)	Annex III screening	30-50K cases/year	Borderline review	EU AI Act conformity
Incident Pattern Analysis	ML + network analysis	Systemic risk detection	10K+ incidents/year	Escalation decision	NIS2 cybersecurity
Standards Gap Analysis	NLP + knowledge graph	Standards prioritisation	1,000+ standards	Work programme decision	ISO/IEC, ETSI
Governance Scenario Modelling	CRISF simulation	Impact analysis	Per-decision	Policy choice	Regulatory impact assessment

3. The AITGF Framework

3.1 Governance Institutions and Deployment Design

Three governance institution deployments over 12 months (October 2024 - September 2025). Institution Alpha (EU member state national AI authority, established under EU AI Act): primary functions: AI system conformity assessment screening, high-risk AI incident investigation; deployed: conformity assessment AI + evidence synthesis + incident analysis; volume: 480 conformity assessment screenings, 84 incident investigations. Institution Beta (European cybersecurity agency, NIS2 competent authority): primary functions: significant cybersecurity incident investigation, critical infrastructure risk assessment; deployed: incident pattern analysis +

evidence synthesis; volume: 2,840 incident reports processed. Institution Gamma (international standards body, AI standards committee): primary functions: AI Act harmonised standard development, standard gap analysis; deployed: standards gap analysis + evidence synthesis + governance scenario modelling; volume: 1,240 technical contributions analysed per standard cycle.

3.2 GDQS Metric and Evaluation Design

$GDQS = 0.35 \times EvidenceQuality + 0.25 \times DecisionAccuracy + 0.25 \times ProcessEfficiency + 0.15 \times DemocraticAccountability$. EvidenceQuality: independent expert assessment of governance decision evidence base quality (vs. matched historical decisions without AI-DSS); criteria: completeness, relevance, accuracy, balance. DecisionAccuracy: for decisions with verifiable outcomes (conformity assessment classification vs. expert panel ground truth, incident severity vs. post-hoc investigation outcome); process verification accuracy. ProcessEfficiency: $1 - (AI_DSS_time / baseline_time)$; baseline = matched historical decisions without AI-DSS. DemocraticAccountability: transparency and auditability of AI-assisted governance decisions. Table 2 presents AITGF specifications.

Table 2: AITGF Framework -- Three Institutions, Five AI-DSS, GDQS Weights

Institution	Primary Function	AI-DSS Deployed	Volume (12 months)	GDQS Priority	Evaluation Benchmark
Alpha (AI Authority)	Conformity assessment + incident	Conformity AI + Evidence + Incident	480 assessments, 84 incidents	Decision Accuracy (0.25)	Expert panel ground truth
Beta (Cyber Agency)	Incident investigation	Incident pattern + Evidence	2,840 incidents	Process Efficiency (0.25)	Post-hoc severity validation
Gamma (Standards Body)	Standard development	Gap analysis + Evidence + Scenario	1,240 contributions/cycle	Evidence Quality (0.35)	Expert standards panel

4. Results

4.1 Regulatory Evidence Synthesis and Conformity Assessment AI

Regulatory evidence synthesis: highest GDQS = 0.908 (consistent with AIDSE-PF evidence synthesis results); Institution Alpha conformity assessment: evidence synthesis processes 1,240 documents per conformity assessment case (vs. 84 manually -- 14.8x coverage); evidence quality rating: 4.4/5.0 vs. 2.8/5.0 historical (+57.1%); Institution Gamma standards work: synthesis of 1,240 technical contributions per standard cycle identifies 28.4% of contributions as directly relevant (vs. 12.4% in manual review -- higher recall at same precision); time to evidence brief: 3.2 days (vs. 18.4 days manual -- 82.6% reduction). Conformity assessment AI: GDQS = 0.876; Annex III classification accuracy: 94.2% vs. expert panel

ground truth (480 cases, 15 borderline cases referred to human review); processing speed: 42 minutes per case (vs. 18.4 hours manual -- 96.2% time reduction for screening stage); false negative rate (missed high-risk classifications): 2.4% -- within acceptable threshold for screening; all false negatives caught at human review stage (two-stage process: AI screening + human review of AI-classified borderline + non-high-risk).

4.2 Incident Analysis, Gap Analysis, Scenario Modelling, and GDQS Summary

Incident pattern analysis: GDQS = 0.884; highest process efficiency gain (68.4% -- incident triage time from 8.4 to 2.4 hours per incident); systemic risk identification: incident pattern analysis detected 3 systemic vulnerability patterns in NIS2 incident data (IoT botnet command-and-control topology, supply chain compromise pattern, AI model poisoning indicator) that were not apparent in individual case review; decision accuracy: incident severity classification 88.4% vs. post-hoc investigation. Standards gap analysis: GDQS = 0.856; identified 42 standards gaps (requirements in EU AI Act Annex IV documentation not covered by existing harmonised standards); standards body work programme prioritised 18/42 gaps as urgent (estimated adoption delay risk > 6 months); time saving: gap analysis 4.2 days vs. 8.4 weeks manual. Governance scenario modelling: GDQS = 0.848; used in 12 regulatory decisions (4 per institution); decision quality improvement vs. no-scenario baseline: +32.4% (expert panel assessment). GDQS: Evidence synthesis 0.908;

Incident analysis 0.884; Conformity AI 0.876; Gap analysis 0.856; Scenario modelling 0.848. Table 3 presents institution GDQS. Table 4 presents dimension GDQS. Figures 1-4 visualise key findings.





Table 3: AITGF GDQS Institution Scores -- Five Architectures x Three Institutions (0-1)

Architecture	Institution Alpha	Institution Beta	Institution Gamma	Overall GDQS
Evidence Synthesis	0.916	0.892	0.916	0.908
Incident Pattern Analysis	0.884	0.900	0.868	0.884
Conformity Assessment AI	0.900	0.852	0.876	0.876
Standards Gap Analysis	0.844	0.836	0.888	0.856
Governance Scenario Model.	0.856	0.836	0.852	0.848

Table 4: AITGF GDQS Dimension Scores -- Five Architectures x Four Dimensions (0-1)

Architecture	Evidence Quality	Decision Accuracy	Process Efficiency	Democratic Accountability	GDQS
Evidence Synthesis	0.960	0.880	0.920	0.840	0.908
Incident Pattern Analysis	0.880	0.884	0.924	0.840	0.884

Archit ecture	Eviden ce Qua lity	Decisi on Acc uracy	Proces s Effic iency	Democ ratic A ccount ab.	GDQS
Confor mity As sessme nt AI	0.880	0.942	0.880	0.840	0.876
Standa rds Gap An alysis	0.880	0.840	0.880	0.760	0.856
Govern ance S cenario Model.	0.880	0.840	0.840	0.800	0.848



5. Discussion

5.1 Conformity Assessment AI at Scale: The 30,000-Case Challenge

The conformity assessment AI result -- 94.2% Annex III classification accuracy at 42 minutes per case (vs. 18.4 hours manual) -- directly addresses the EU AI Act's implementation scaling challenge. The European Commission estimates 30,000-50,000 AI systems will require conformity assessment annually under the EU AI Act; national AI authorities are currently staffed for hundreds of cases per year at most. Without AI-assisted conformity assessment, the EU AI Act faces an implementation credibility crisis: requirements that cannot be enforced at scale become de facto voluntary. The two-stage architecture (AI screening + human review of borderline cases) is the appropriate governance design: AI handles the

97.6% of cases that are clearly high-risk or clearly not high-risk; human experts focus on the 2.4% borderline cases that require judgment. The 2.4% false negative rate (missed high-risk classifications) is acceptable because all AI-classified non-high-risk cases undergo human sampling review (random 5% of AI-classified non-high-risk reviewed by human expert as quality control). Continuous improvement loop: every human-reviewed borderline case is added to the training dataset; quarterly model retraining incorporates novel AI system types and evolving Annex III precedents; accuracy is expected to increase toward 97-98% as the training dataset expands.

5.2 AI-DSS and the Human-AI Governance Partnership

AITGF deployments consistently confirm the HAIRAF advisory AI recommendation (Horvath and Petrov, 2025): AI-DSS achieves highest governance quality when it functions as an advisory layer that augments rather than replaces human governance judgment. The evidence synthesis and conformity assessment AI deployments both use a strict advisory design: AI outputs are presented as "evidence brief" and "classification proposal" not as "decision recommendation" -- the framing preserves human accountability and prevents the automation bias that would degrade governance quality if AI recommendations were treated as authoritative. The democratic accountability dimension scores (0.800-0.840) are the weakest GDQS dimension

across all architectures -- reflecting that transparency and citizen-facing accountability of AI-assisted governance is the least developed aspect of current deployments. AITGF recommends mandatory public disclosure of AI-DSS deployment in all regulatory governance contexts (what AI systems assisted which decisions), consistent with the PACE-TGF democratic accountability dimension (Novak, 2025).

6. Conclusion

6.1 Summary

AITGF evaluated five AI-DSS architectures across three governance institutions over 12 months. Key results: evidence synthesis GDQS = 0.908 (14.8x document coverage vs. manual); incident pattern analysis process efficiency gain 68.4%; conformity assessment AI 94.2% Annex III accuracy at 96.2% time reduction; standards gap analysis identifies 42 EU AI Act gaps in 4.2 days vs. 8.4 weeks; all architectures GDQS > 0.840. Democratic accountability is the weakest dimension -- mandatory AI-DSS disclosure recommended.

6.2 Open Resources

The AITGF deployment guide, GDQS calculator, conformity assessment AI model (HuggingFace: `aitgf/euaiact-annex3-bert`), evidence synthesis pipeline (regulatory domain), incident pattern analysis toolkit, standards gap analysis code (knowledge graph + NLP), governance scenario modelling API (CRISF integration), and 12-month deployment dataset are available at

aitgf-governance.org under Apache 2.0 License.

References

- Bianchi, L., and Garcia, E. (2024). Ethics-by-design models for socio-technical systems. *Journal of Ethics in Emerging Technologies & Society*, 2024(1), 9-16.
- 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).
- European Commission (2022). NIS2 Directive. Directive (EU) 2022/2555.
- Gabriel, I. (2020). Artificial intelligence, values, and alignment. *Minds and Machines*, 30(3), 411-437.
- Horvath, N., and Petrov, H. (2025). Human-AI interaction models for socially responsible automation. *Journal of Ethics in Emerging Technologies & Society*, 2025(1), 17-26.
- Horvath, N. (2025). Computational simulation of technology regulation and social outcomes. *Journal of Ethics in Emerging Technologies & Society*, 2025(3), 39-47.
- 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 governing emerging technologies. *Journal of Ethics in Emerging Technologies & Society*, 2025(3), 29-38.
- OECD (2023). Regulatory AI: Prospects and Challenges. OECD Digital Economy Papers.
- Petrov, N., and Klein, C. (2025). Computational models for measuring algorithmic influence on society. *Journal of Ethics in Emerging Technologies & Society*, 2025(1), 1-8.
- 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.
- Shneiderman, B. (2022). *Human-Centered AI*. Oxford University Press.
- Sunstein, C. R. (2017). *Human Agency and Behavioral Economics*. Palgrave Macmillan.
- Vallor, S. (2016). *Technology and the Virtues*. Oxford University Press.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.

Declarations

Funding

This research was supported by the Italian Ministry of University and Research (MUR) under grant PRIN2025-2027-AITGF (AI-DSS for Technology Governance Framework) and the Swiss National Science Foundation (SNSF) under grant 200021-230024. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper. The three governance institutions participated under research partnership agreements that did not restrict publication.

Data Availability Statement

The AITGF deployment guide, GDQS calculator, conformity assessment AI model, evidence synthesis pipeline, incident analysis toolkit, gap analysis code, scenario modelling API, and deployment dataset are available at aitgf-governance.org under Apache 2.0 License.

Ethical Approval

This study evaluates governance decision support processes under institutional research partnership agreements. No personal data was accessed. Ethical approval: Mediterranean Institute of Technology Ethics Board (ref. MITR-2025-ETH-0184).

Appendix A

AITGF System Specifications and GDQS Calculation

Technical specifications and GDQS metric calculation
for AITGF AI-DSS architectures.

Conformity Assessment AI Specification

GDQS Calculation Protocol