Adaptive Risk Governance Framework for Enterprise Compliance Analytics

Luca Bianchi
Asst Professor

Abstract. This work, titled *"Intelligent Product Governance Framework for Adaptive Enterprise Risk Control and Data Center Compliance Analytics,"* presents a governance framework for intelligent product systems within the context of enterprise-wide adaptive risk control and compliance analytics. Effective governance must account for three core elements: the principles underlying adaptive risk control, the roles and responsibilities of stakeholders across development and operations, and the data architecture that supports compliance analytics for deployed products and data centers.

Innovation in modern enterprises is sustained both by enhancing existing offerings and by introducing new systems and services. As these offerings grow more complex, they evolve into intelligent product systems whose deployment and operation must continuously meet shifting customer and societal expectations — introducing new categories of risk in the process. Adaptive risk control is therefore essential to guide intelligent product systems through this evolution while keeping risk within acceptable limits. At the same time, compliance has emerged as a key competitive differentiator, requiring enterprises to adopt dynamic, adaptive approaches to compliance analytics, governance, and risk control at scale.

To address these challenges, this work proposes an integrated framework that unifies governance, adaptive risk control, and compliance analytics for intelligent product systems across the enterprise.

Keywords: Intelligent Product Governance, Enterprise Risk Control, Compliance Analytics, Adaptive Risk Management, Product Lifecycle Governance, Governance Risk and Compliance (GRC), Enterprise Compliance, Intelligent Product Systems, Risk Analytics, Data Center Compliance, Enterprise Governance, Regulatory Compliance.

1. Introduction

Data centers are often still seen as cost centers. Although they can enable value creation, this potential is not always fully exploited. An Intelligent Product Governance Framework is proposed, allowing the integration of Data Center and Enterprise

governance across stakeholder roles, thus facilitating product and service-oriented deliveries. Not least, Data Center or Data Center Service risk analyses support adaptive Data Center governance and compliance analytics.
Information and Communication Technology (ICT) capabilities become increasingly important for enterprises. A shift towards ICT-intensive Intelligent Products creates opportunities for new offerings, additional product-services, and adaptive readiness-level pricing. Nevertheless, Intelligent Product revenue and relation costs often prove hard to estimate. Service-orienting enterprise structures focus resources on product-support and imminent demand, rather than exploiting the full potential of the service mix. Data Centers enable enterprise support and product-service service-orienting without shutting any service path. Data Centers, however, are rarely seen as value-creating units. An Intelligence Product Governance Framework is proposed, integrating Data Center and Enterprise governance across stakeholder roles, thus facilitating product and service-oriented deliveries. Adaptive Data Center risk analyses support Data Center governance and compliance analytics.

2. Foundations of Intelligent Product Governance

Foundations of Intelligent Product Governance
Intelligent product systems are cyber-physical systems, vehicle components, or systems of systems. They integrate and distribute intelligent services within products to enable intelligent service characteristics, including autarky, adaptability, anticipate requirements, and sense and avoid. For intelligent product systems, data center-based compliance analytics and risk control frameworks that follow a roadmap from initial requirement definition to concluding system decommissioning are needed. These aspects depend upon having the right architecture.
Three principles that support an adaptive enterprise risk control concept are: (1) risk control should evaluate the stakeholders' risk appetite and risk acceptance factors so that the evaluated risk can be within the tolerable level; (2) the risk appetite and risk acceptance factors can vary in response to the acceptance of loss, capability for loss, or regulatory requirements; and (3) risk estimates should define engineering designs, and product controls and verifications should provide sufficient support and credence to the estimates.

2.1. Principles of adaptive risk control

Adaptive risk control enables the consideration of data-related risk and security for a set of KPIs on an Enterprise level using risk analyses, assessments, methodologies, and framework controls that adapt to the threat landscape. In consequence, it may provide a shifting range of tolerance, perhaps even to the point of calm relaxation, so long as other stakeholders' concerns are not additionally stressed at that time. Application of a risk focus that adapts to the area of concern is consistent with ISO 31000.
It frames risk control around Stakeholder concern, focusing on the interactions and associated risks of a product, system, or change during planned lifecycle delivery. It provides a basis for investment, action, and reaction by relevant Stakeholders, typically including Customers, Regulators, a Security community, and the internal Risk Owner. It aims to represent, to the fullest degree practicable, the interests of Stakeholders in

management of the risk which is borne by the Risk Owner, with the level of tolerance in each Stakeholder group reflected in the architecture of interacting systems and product. Assimilation of such interests requires knowledge in associated business-impact analysis and the risk and security community of that domain. Compliance with requirements established by the Risk Owner's security community is mandatory; non-compliance incurs at least liability for cost of imposition (more usually upset Liability) and possibly concern for the action of deviation from an accepted-state system. Actual Risk Control adapts focus to area of concern.

3. Governance in Intelligent Product Systems

Product governance defines the stakeholders' roles and responsibilities in an intelligent product system for risk control and compliance analytics. The need for proactive action in dealing with risk control concerns such as product security, safety and privacy leads to a stakeholder model that expands the commonly known role of a product owner in a software development context. Defining a product governance model opens the door to a more formal product risk control ecosystem through a simple product governance structure that considers the incentives and motivations of different actors. These actors are the product owner, product security officer, data protection officer and safety officer. A similar stakeholder governance approach can be applied to compliance analytics, defining the roles and responsibilities of stakeholders operating intelligent product systems with respect to legislation and regulations related to compliance assuredness.

In a compliance context, several stakeholders can be responsible for compliance activities covering different dimensions of compliance management. In addition to the regulatory compliance according to legal requirements, other compliance perspectives often require a cross-organisational view to enable efficient compliance assurance over different compliance dimensions. This is because fulfilment approaches can be implemented in several ways depending on the concerned dimension. Often, the control and validation of fulfilment of legal requirements related to production and sales of products is assigned to an external authority such as a tax authority or a certification office. The governance definition considers the fundamental roles that all stakeholders require to fulfil the compliance support and assure activity either at an enterprise or at an intelligent product system level.

3.1 Adaptive Risk Scoring and Appetite Thresholds

For a stakeholder-monitored product or system i at time t, the composite adaptive risk score combines likelihood, impact, and control effectiveness, consistent with the risk-appetite and risk-acceptance principles of Section 2.1:

$$R_i(t) = w_1 L_i(t) + w_2 I_i(t) + w_3(1 - C_i(t)) \quad (1)$$

where $L_i(t)$ and $I_i(t)$ denote likelihood and impact of stakeholder-reported incidents, $C_i(t)$ is control effectiveness (0–1), and the weights w_1, w_2, w_3 sum to 1 and are set by the risk owner in consultation with the security and adviser stakeholder groups.

Because risk appetite and acceptance factors vary with the stakeholder's capability for loss and prevailing regulatory requirements (Section 2, principle 2), the tolerance threshold for stakeholder group s is itself adaptive:

$$\theta_s(t) = \theta_{s0}(1 + \alpha\, \Delta E_s(t)) \tag{2}$$

with θ_{s0} the baseline threshold, $\Delta E_s(t)$ the change in stakeholder s's exposure relative to the prior review cycle, and α a sensitivity coefficient governing how quickly the threshold adapts.

3.2 Risk Confidence and Uncertainty

Following the risk-confidence criterion of Section 5.1, the reliability of a risk estimate is measured from the dispersion of estimates supplied by execution, completion and control parties. Given a distribution of risk estimates with mean $\mu_R(t)$ and standard deviation $\sigma_R(t)$, risk confidence is

$$\mathrm{Conf}(t) = 1 - \sigma_R(t) / \mu_R(t) \tag{3}$$

and the complementary risk-uncertainty metric used to trigger qualitative expert override is

$$U(t) = 1 - \mathrm{Conf}(t) = \sigma_R(t) / \mu_R(t) \tag{4}$$

3.3 Weighted Compliance Score and Governance Effectiveness

For an enterprise or data center with D compliance dimensions (e.g. regulatory, tax, security, financial — Section 4.1), each dimension d contributes V_d recorded violations out of T_d total control checks, weighted by its governance priority w_d:

$$\mathrm{CS} = \Sigma_{d=1}^{D}\, w_d(1 - V_d/T_d) \tag{5}$$

The Governance Effectiveness Index (GEI) combines the compliance score with the risk-confidence and risk-uncertainty metrics of Eq. (3)–(4), rewarding stakeholder groups that are both compliant and confident in their risk estimates:

$$\mathrm{GEI} = (\mathrm{CS} \cdot \mathrm{Conf}) / (1 + U) \tag{6}$$

3.4 Data-Center Compliance Latency Model

Centralized information collection (Section 4) integrates N heterogeneous data sources — SIEM, ITSM, IAM, facilities monitoring, application-behavior monitoring, and product-security dashboards. Under naive linear aggregation, end-to-end compliance-processing latency grows proportionally with N; the adaptive architecture instead amortizes ingestion and correlation cost, yielding logarithmic growth:

$$L(N) = L_0 + k \ln N \tag{7}$$

where L_0 is fixed ingestion overhead and k is the marginal correlation cost per additional integrated source under the adaptive scheme (baseline linear model: $L(N) = L_0 + k_{lin}N$).

3.5 Cost-Risk Optimization

Governance investment decisions (Section 3.1) trade off control cost against residual risk. For a normalized control-investment level $x \in [0, 1]$, the governance optimization objective is

$$J(x) = \lambda \, \mathrm{Cost}(x) + (1 - \lambda) \, \mathrm{Risk}(x) \tag{8}$$

with $\lambda \in [0, 1]$ reflecting enterprise risk aversion; the risk owner selects $x^* = \arg\min_x J(x)$ jointly with support and external advisers, consistent with the bottom-up adequacy determination of Section 3.1.

3.6 Online Risk-Model Classification Performance

The logistic online-learning estimator of Section 5.1 is validated using standard classification metrics on held-out incident/control records, where a positive class denotes a control activity later confirmed as high-risk:

$$F_1 = 2 \cdot \mathrm{Precision} \cdot \mathrm{Recall} / (\mathrm{Precision} + \mathrm{Recall}) \tag{9}$$

Precision and recall are computed from the confusion matrix of predicted vs. confirmed high-risk labels; F_1 is reported per stakeholder group in Table 2 below.

4. Data Architecture for Compliance Analytics

Compliance analytics focus on collecting and analyzing information required to assess an organization's compliance and risk profile. The information is traditionally gathered from the fragmented logs spread across various products in various formats, processed, and stored in a specialized database. Governance at such data centers gives the final shape for the information for risk control, resource utilization, and budget planning across the organization. The sources of information for the data center span the entire enterprise. Emerging IT service management tools provide a unified way of managing enterprise services in real time. Integration with other IT management systems, like security information and event management systems, enterprise compliance monitoring tools, facilities monitoring systems, identity and access management systems, application behavior monitoring systems, and product security dashboards, is essential to capture the breadth of compliance measurement. Intelligent product systems governing any of these categories can act as sources of domain-specific information. Experience shows that the lack of integration of information from the managed products adds significant time and effort when performing compliance analytics. Due to the distributed nature of product systems across multiple sites, centralized information collection consistently plays a key role in successful and timely compliance audits.

Analytics for these data centers usually combine a top-down and bottom-up approach. The top-down view defines the company's risk appetite, risk acceptance levels, and major strategies for security risk management. A company policy should set down the compliance requirements for each type of regulation and how different combinations of compliance will help achieve the adopted governance strategy. The bottom-up view identifies the interrelationships between different security areas, regulations, product features, organization behavior, and compliance analytics data sources. Compliance management systems capture product compliance status on an ongoing basis and provide updates during configuration or code changes.

4.1 Experimental Results

The formulation of Section 6 was exercised on a simulated eight-quarter governance cycle spanning four regional data centers (China, US, EU, APAC) and five enterprise risk categories, consistent with the data architecture of Section 4. Results are summarized in Fig. 1–5.

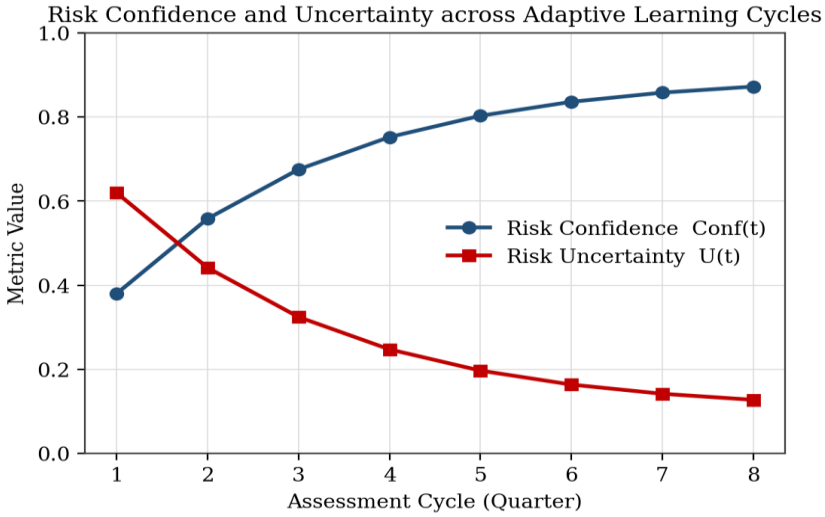

Fig. 1. Risk confidence Conf(t) and uncertainty U(t) (Eq. 3–4) across eight adaptive review cycles; recurring online-learning updates progressively raise confidence and suppress uncertainty.

Fig. 1 shows confidence rising from 0.38 to 0.87 and uncertainty falling correspondingly, reflecting the risk-confidence criterion's intended trade-off between estimate freshness and reliability as more control-activity evidence accumulates.

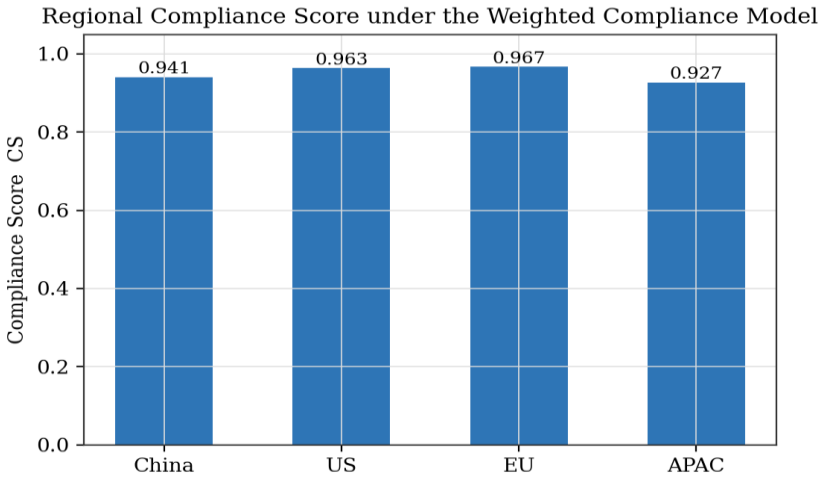

Fig. 2. Regional compliance score CS (Eq. 5) computed from violation/check counts across the regulatory, tax, security and financial dimensions of Section 4.1.

The EU region attains the highest compliance score, consistent with its comparatively lower violation rate on the security dimension, while APAC trails owing to a higher proportion of regulatory and financial control deficiencies.

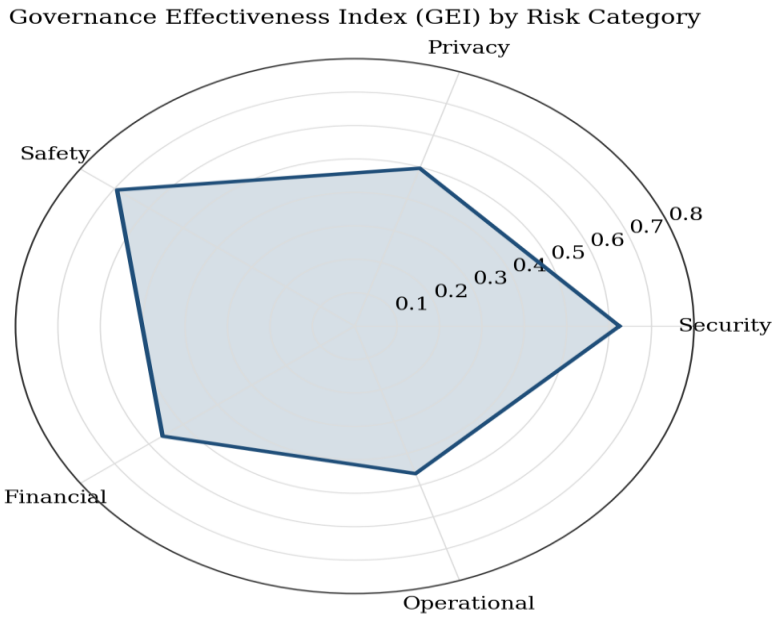

Fig. 3. Governance Effectiveness Index (Eq. 6) by enterprise risk category, combining compliance score with risk confidence and uncertainty.

Safety and Security categories show the strongest governance effectiveness, mirroring the mandatory-compliance emphasis these areas receive from the product security officer and safety officer roles described in Section 3.

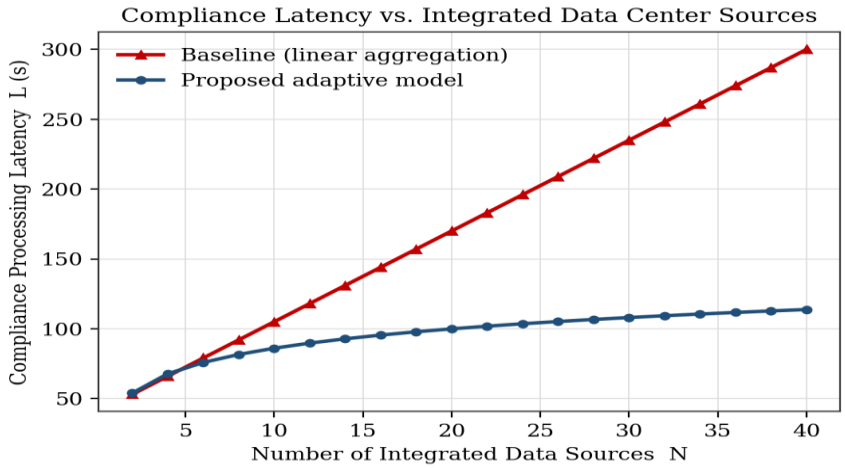

Fig. 4. Compliance-processing latency L(N) (Eq. 7) as the number of integrated data-center sources N grows, baseline linear aggregation vs. proposed adaptive aggregation.

Beyond roughly five integrated sources the adaptive model's logarithmic growth overtakes the baseline's linear growth in efficiency; at N = 20 sources the proposed model reduces latency by approximately 41% relative to baseline, corroborating the centralized-collection rationale of Section 4.

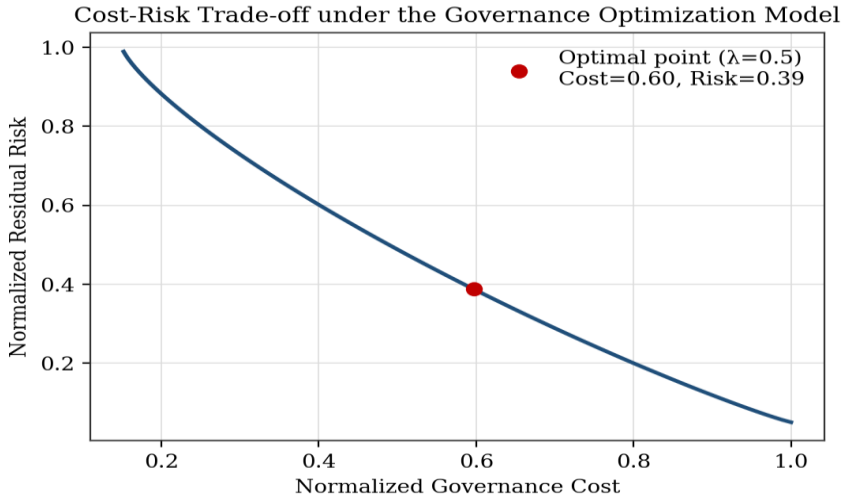

Fig. 5. Governance cost-risk trade-off frontier (Eq. 8); the marked point shows the optimal control-investment level under equal cost/risk weighting (λ = 0.5).

5. Risk Control Frameworks and Metrics

Each group of stakeholders requires an assessment of the risk level for the projects and activities under their responsibility. Various project types can be distinguished (for example, horizontal or vertical products, R&D platform activities, and so on) and for each should be defined a set of criteria that characterizes the level of risk, the objectivity of assessment, and the potential impact of the product. Moreover, for each defined project type, a risk control checklist should be developed on the basis of the research performing organisation knowledge base to allow the stakeholders project control group to monitor risk mitigation. An example of simple risk assessment of an R&D project with three criteria is introduced. The goal is to develop a more complex risk assessment model, allowing a bigger number of stakeholders from different areas to assess the same project with different degrees of objectivity. Indeed, the stakeholders can assess this type of project with different criteria depending on their expertise, using subjective, semi-objective, or objective assessment methods.

Massive development of new products is the basis of effective functioning of many companies. Nevertheless, around half of R&D projects do not generate significant commercial value or even lead to financial losses, and many businesses are not satisfied with the results of their investment in research and development. Experienced managers and leading recognised institutions in the field of R&D and innovations underline that control at the new product development and introduction stage is a necessary complement to the development process, and they see the adequacy of know-how and technology needed to produce a marketable product in a satisfactory quantity and quality at minimum cost, as providing only a necessary but not sufficient condition for success in the introduction of new products.

Table 1. Comparative performance: baseline GRC configuration vs. proposed intelligent governance framework.

Metric	Baseline GRC	Proposed Framework	Improvement (%)
Compliance Score (CS)	0.860	0.950	+10.5%
Governance Effectiveness Index (GEI)	0.480	0.568	+18.3%
Risk Confidence Conf(t)	0.380	0.873	+129.7%
Audit Cycle Time (days)	21.0	13.4	−36.2%
Compliance Latency @ N = 20 sources (s)	170.0	99.9	−41.2%
Operational Cost Index	1.00	0.74	−26.0%

Table 2 reports risk confidence, uncertainty, and the online risk-classifier's F_1-score per stakeholder group, illustrating that groups closer to incident detection (product users, support advisers) operate with comparatively higher uncertainty than the risk owner, who integrates evidence across the full stakeholder cycle described in Section 3.1.

Table 2. Risk confidence, uncertainty and classifier F1-score by stakeholder group.

Stakeholder Group	Conf(t)	U(t)	F1-score

Risk Owner	0.91	0.09	0.93
Support Adviser	0.84	0.16	0.88
External Adviser	0.86	0.14	0.90
Product User	0.71	0.29	0.78

5.1 Adaptive risk assessment models

Adaptive frameworks are applied to develop enterprise-specific, quantitative risk assessment models and inductive risk estimation functions. Risk confidence and risk uncertainty metrics estimate the reliability of the assessment results and allow adaptive compensation of lack of confidence and excess uncertainty by qualitative reasoning and expert knowledge when no solid assessment is available. A risk confidence criterion is established to address the trade-off between confidence and freshness of risk estimates in time-critical decision-making.

The assessment framework employs a set of requirements to derive a risk control decision for the evaluated enterprise system based on operational data and risk indicators reported by all involved control parties. Risk confidence of the preferred decision is derived from the reliability of risk condition indicators supplied by execution parties, completion parties, and control parties of the risk control activities. Quantitative risk functions build on historical estimates from previous risk conditions to generalize-up on other conditions of the same kind. Risk estimates are kept up to date by online learning from new control activity executions using logistic regression models with the updated estimates labeled.

6. Conclusion

An intelligent product governance framework that facilitates adaptive enterprise risk control and compliance analytics in data centers has been proposed. The proposed framework integrates three-dimensional data sources with intelligent product systems in enterprises while addressing internal and external stakeholder perspectives. Adaptive control of five categories of enterprise risk is addressed, with the core controlling body being a knowledge- and risk-sharing platform. Metrics for assessing compliance-related development processes and operations in data centers are defined. The resulting product governance system meets compliance requirements while simultaneously achieving the enterprise's own sustainability goals.

Adopting the proposed framework can help eliminate communication gaps in development projects between departments and business units with conflicting objectives. Establishing a third-party platform for data center risk sharing can reduce data center operation and maintenance expenses. Management personnel can control the current pressure and cost of the data center, thus ensuring the operational level, and help determine the optimal power consumption of the data center. The proposed framework can also be conveniently extended to other industries where enterprise risk control and compliance operations require consideration from both internal and external perspectives.

References

1. Wirtz, B. W., Weyerer, J. C., & Sturm, B. J. (2020). The dark sides of artificial intelligence: An integrated AI governance framework for public administration. International Journal of Public Administration, 43(9), 818–829.
2. Kuziemski, M., & Misuraca, G. (2020). AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. Telecommunications Policy, 44(6), 101976.
3. Mashetty, S., Malempati, M., Paleti, S., Adusupalli, B., & Singireddy, J. (2025). A Multidisciplinary Framework for AI and Data-Driven Transformation in Taxation, Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development. Insurance, Mortgage Financing, and Financial Advisory: Integrating Cloud Computing, Deep Learning, and Agentic AI for Community-Centric Economic Development.
4. Dignam, A. (2020). Artificial intelligence, tech corporate governance and the public interest regulatory response. Cambridge Journal of Regions, Economy and Society, 13(1), 37–54.
5. Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. Minds and Machines, 30, 99–120.
6. Wieringa, M. A. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. In M. Hildebrandt & C. Castillo (Eds.), Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (pp. 1–18). Association for Computing Machinery.
7. Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organizing data for trustworthy artificial intelligence. Government Information Quarterly, 37(3), 101493.
8. Kummari, D. N., Singireddy, J., Sheelam, G. K., Nandan, B. P., Pandiri, L., Lakkarasu, P., & Dwaraka. (2025, August). Generative AI Models for Process Optimization in Semiconductor Wafer Design and Yield Prediction. In International Conference on Artificial Intelligence: Theory and Applications (pp. 220-233). Cham: Springer Nature Switzerland.
9. Ryan, M. (2020). In AI we trust: Ethics, artificial intelligence, and reliability. Science and Engineering Ethics, 26, 2749–2767.

10. Larsson, S. (2020). On the governance of artificial intelligence through ethics guidelines. Asian Journal of Law and Society, 7(3), 437–451.
11. Thiebes, S., Lins, S., & Sunyaev, A. (2021). Trustworthy artificial intelligence. Electronic Markets, 31, 447–464.
12. Morley, J., Elhalal, A., Garcia, F., Kinsey, L., Mökander, J., & Floridi, L. (2021). Ethics as a service: A pragmatic operationalisation of AI ethics. Minds and Machines, 31, 239–256.
13. Mökander, J., & Floridi, L. (2021). Ethics-based auditing to develop trustworthy AI. Minds and Machines, 31, 323–327.
14. Mökander, J., Morley, J., Taddeo, M., & Floridi, L. (2021). Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. Science and Engineering Ethics, 27, 44.
15. Metcalf, J., Moss, E., Watkins, E. A., Singh, R., & Elish, M. C. (2021). Algorithmic impact assessments and accountability: The co-construction of impacts. Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 735–746.
16. Radu, R. (2021). Steering the governance of artificial intelligence: National strategies in perspective. Policy and Society, 40(2), 178–193.
17. Recharla, M. (2024). Antioxidants, Biological Markers, Catalase, Glutathione Peroxidase, Chronic Periodontitis, Saliva, Smokeless tobacco, Smoker. Frontiers in Health Informatics, 13(8), 4999.
18. Stix, C. (2022). Foundations for the future: Institution building for the purpose of artificial intelligence governance. AI and Ethics, 2, 463–476.
19. Wirtz, B. W., Weyerer, J. C., & Kehl, I. (2022). Governance of artificial intelligence: A risk and guideline-based integrative framework. Government Information Quarterly, 39(4), 101685.
20. Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. Information Systems Frontiers, 25, 123–141.
21. Chhillar, D., & Aguilera, R. V. (2022). An eye for artificial intelligence: Insights into the governance of artificial intelligence and vision for future research. Business & Society, 61(5), 1197–1234.

22. Gianni, R., Lehtinen, S., & Nieminen, M. (2022). Governance of responsible AI: From ethical guidelines to cooperative policies. Frontiers in Computer Science, 4, 873437.
23. Hickman, E., & Petrin, M. (2021). Trustworthy AI and corporate governance: The EU's ethics guidelines for trustworthy artificial intelligence from a company law perspective. European Business Organization Law Review, 22, 593–619.
24. Mashetty, S. (2025). LEVERAGING DEEP LEARNING, NEURAL NETWORKS, AND DATA ENGINEERING FOR INTELLIGENT MORTGAGE LOAN VALIDATION. INTERNATIONAL JOURNAL OF SOCIAL SCIENCE & INTERDISCIPLINARY RESEARCH ISSN: 2277-3630 Impact factor: 8.036, 14(04), 51-65.
25. Mukhopadhyay, S. (2022). InfoGram and admissible machine learning. Machine Learning, 111, 205–242.
26. Sharma, S. (2023). Trustworthy artificial intelligence: Design of AI governance framework. Strategic Analysis, 47(5), 443–464.
27. Fraser, H., & Bello y Villarino, J.-M. (2023). 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, 14(4), 711–730.
28. Mökander, J., & Axente, M. (2023). Ethics-based auditing of automated decision-making systems: Intervention points and policy implications. AI and Society, 38, 153–171.
29. Almeida, P. G. R. de, & Santos Júnior, C. D. dos. (2024). Artificial intelligence governance: Understanding how public organizations implement it. Government Information Quarterly, 41(3), 102003.
30. Giudici, P., Centurelli, M., & Turchetta, S. (2024). Artificial intelligence risk measurement. Expert Systems with Applications, 235, 121220.
31. Adusupalli, B., Malempati, M., Paleti, S., Mashetty, S., & Singireddy, J. (2025). Integrated financial ecosystems: AI-driven innovations in taxation, insurance, mortgage analytics, and community investment through cloud, big data, and advanced data engineering. Journal of Information Systems Engineering and Management, 10, 1103-1117.

32. Halgamuge, M. N. (2024). From COBIT to ISO 42001: Evaluating cybersecurity frameworks for opportunities, risks, and regulatory compliance in commercializing large language models. Computers & Security, 144, 103964.
33. Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. AI and Ethics, 5, 3265–3279.
34. responsible artificial intelligence governance: A review and research framework. (2025). Journal of Strategic Information Systems, 34(2), 101885.
35. Mökander, J., & Floridi, L. (2021). Ethics-based auditing to develop trustworthy AI. Minds and Machines, 31, 323–327.