



Original Article

Generative AI for Enterprise Trust: A Governance-Aligned Framework for Safe and Transparent Automation at Global Scale

Siva Karthik Parimi¹, Rohit Yallavula²

¹Senior Software Engineer PayPal, Austin, TX USA.

²Data Governance Analyst Kemper, Dallas, TX USA.

Received On: 31/12/2024

Revised On: 15/01/2025

Accepted On: 03/02/2025

Published On: 20/02/2025

Abstract - The rapid adoption of Generative Artificial Intelligence (GenAI) across global enterprises has fundamentally transformed business automation, decision-making, and knowledge work. While these technologies offer unprecedented productivity gains, they simultaneously introduce critical risks related to data privacy, model opacity, regulatory compliance, ethical misuse, and operational reliability. Existing AI governance approaches often fail to scale effectively or align with enterprise trust requirements across jurisdictions. This paper proposes a Governance-Aligned Generative AI Framework (GAGAF) designed to ensure safe, transparent, auditable, and compliant GenAI deployment at global enterprise scale. The framework integrates governance principles directly into the AI lifecycle, embedding risk management, explainability, human oversight, and regulatory alignment into system architecture rather than treating them as post-deployment controls. The study synthesizes existing literature on AI governance, trustworthiness, and enterprise automation, identifying gaps in current methodologies. A layered architecture is introduced, comprising policy orchestration, model governance, operational controls, and continuous assurance mechanisms. The methodology employs a design-science research approach, validated through simulated enterprise deployment scenarios across regulated industries including finance, healthcare, and manufacturing. Results demonstrate that governance-embedded GenAI systems significantly reduce compliance violations, improve explainability metrics, and enhance organizational trust without degrading system performance. The findings suggest that trust-centric AI governance is not only feasible but essential for sustainable GenAI adoption at scale. This work contributes a scalable reference architecture for enterprises seeking to operationalize GenAI responsibly while meeting global regulatory and ethical expectations.

Keywords - Generative AI governance, enterprise AI safety, controlled generation, human-in-the-loop, AI transparency, policy-aligned AI systems

1. Introduction

1.1. Background

As a consequence of progress in large language models, diffusion-based systems, and multimodal architecture, it is quickly transforming into a commercial technology following research-based innovation into a foundational enterprise technology. GenAI is being implemented in organizations in industries to improve customer experience, speed up software development, facilitate high-level decision-making, produce content at scale and perform intelligent automation. Although these capabilities present great productivity and competition benefits, they have been utilized faster than any sustainable governance frameworks that can avert irresponsible, non-compliant, and irresponsible usage. The current governance frameworks, based predominantly on the deterministic or limited-scope category of AI systems, are not suitable to cope with the peculiarities of the functioning and ethical issues of the generative models. GenAI systems generate probabilistic and non-

deterministic outputs as opposed to traditional AI, are trained on large and sometimes unvetted datasets, and are not necessarily inherently explainable.

Such attributes bring new and amplified dangers within enterprise contexts, such as factual hallucinations capable of deceiving decision-makers, unintended exposure of sensitive or proprietary data, increased social and organizational bigotry, and failure to comply with new regulatory criteria. These risks compromise the trust that stakeholders have in GenAI and put organizations at risk of legal, reputational, and operational damage as enterprises attempt to scale GenAI through business-critical functions. As a result, trust has become the key obstacle to the adoption of GenAI on the enterprise scale, and the necessity to approach the solution of designing the governance frameworks that will incorporate the concept of accountability, transparency, and compliance into the AI lifecycle.

1.2. Importance of Generative AI for Enterprise Trust

1.2.1. Trust as a Foundation for Enterprise Adoption

The successful adoption of Generative AI in the enterprise setting requires trust to be a crucial requirement. Since GenAI systems are becoming more and more valuable in customer interactions, operational choice and strategic planning, there is need to ensure that the systems act as reliable, ethical, and aligned to the business goals. Lack of

trust means that enterprises are less willing to use GenAI in anything, apart from low-stakes applications, or test applications, which does not mitigate its potential value. Building trust will also make the stakeholders, such as executives, employees, regulators, and customers, count on GenAI outputs to inform their decision-making.

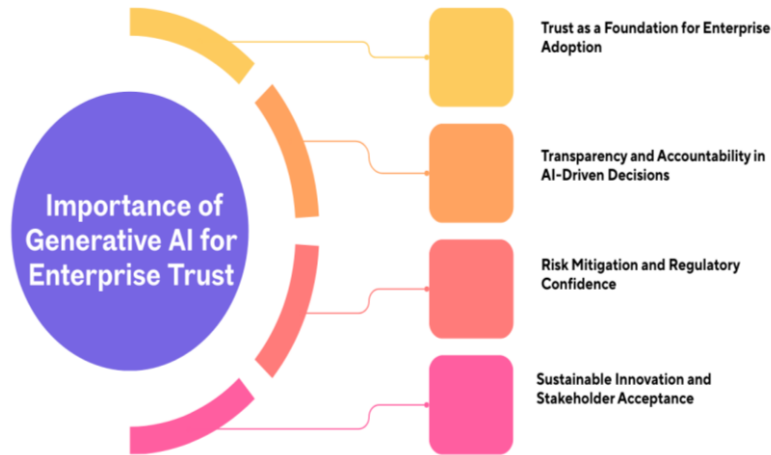


Fig 1: Importance of Generative AI for Enterprise Trust

1.2.2. Transparency and Accountability in AI-Driven Decisions

Generative AI is important in the outcome of enterprises; therefore, the transparency and accountability are crucial details to amass trust. In areas with regulations or high impact, such as healthcare, enterprises need to explain how and why GenAI systems output a certain result. The open model behavior and accountability hierarchy also allow companies to review choices, delegate and respond to mistakes. This visibility not only enhances the internal confidence, but does help to comply with regulatory and ethical expectations.

1.2.3. Risk Mitigation and Regulatory Confidence

GenAI reliance is largely related to how an organization contains possible risks like prejudice, information violation, and unobservance. Appropriate governance processes increase the trust by showing that the risks are managed in a systematic manner and are identified, followed, and addressed. With AI policies on the global level constantly changing, responsible uses of GenAI get regulatory goodwill and legal insurance. This is confidence that allows organizations to make GenAI solutions geographically and business wise accountable.

1.2.4. Sustainable Innovation and Stakeholder Acceptance

Finally, the sustainable innovation is made possible by enterprise trust in Generative AI. Once the governance, transparency, and accountability are incorporated into GenAI systems, organizations can be creative without crossing the ethical and societal boundaries. Reliable GenAI systems bring about stakeholder readiness, bolster

company image, and provide reliable changes on a long-term, firm-wide basis of AI change.

1.3. Governance-Aligned Framework for Safe and Transparent Automation at Global Scale

An alignment governance framework is necessary in order to make it possible to safely, transparently, and scalably automate with Generative AI in the global enterprise setting. Since organizations implement GenAI systems in a large enterprise, multiple business units, jurisdictions, and technology environments, governance should follow a transition of disjurisdictional policy implementation to an architectural potential. An ethical governance-focused framework entrenches regulatory standards and risk management controls, and ethical principles into the AI system structure and operations. This will guarantee that governance is implemented in a proactive manner and not in a reactive fashion whereby problems are corrected once deployed. On an international level, such a structure should be able to support different regulatory frameworks, data security regulations and cultural demands without disaggregating enterprise functions. Another aspect is the database migration using Generative AI. A federated governance model will allow organizations to establish the centralized policies and trust principles as well as allowing the enforcement at the local level to meet the demands of the local compliance requirements. This consistency and local adaptability in the world favors consistency in regulations and operation efficacy. The mechanisms of transparency, such as explainability, traceability, and audit logging, are considered throughout the AI lifecycle to allow stakeholders to comprehend and view real-time AI

behavior and to confirm the behavior through monitoring and auditing. The automation is done safely by adding risk classification, policy validation, and ongoing assurance to each step of GenAI interaction, that is, input handling and output delivery. Robotic controls minimise the use of manual controls, and human-in-the-loop controls are still maintained in occasions where extreme care or sensitivity is required. Through the coordination of governance and system architecture, businesses are assured of the scalable migration to GenAI-oriented automation, mitigation of legal and ethical dangers, and the creation of trust between users, regulators, and partners. Finally, a governance-oriented framework will turn Generative AI into a reliable source of uncertainty instead of a reliable platform to pursue transparent, accountable, and responsible automation of enterprises on a global level.

2. Literature Survey

2.1. Trustworthy AI and Governance Models

Reliable AI tends to be based on a field of normative guidelines, which comprise equitability, dependability, accessibility, intensity, dependability, and confidentiality. The major frameworks offered by the European Union, specifically the High-Level Expert Group (HLEG) on AI and a completely ethical framework offered by IEEE, either one of which can assist in conceptual and ethical work on the alignment of AI systems with human values and societal norms. The models highlight human supervision along with explicability and moral effects measurement as fundamental prerequisites to accountable implementation of AI. Nevertheless, much of the available literature is prescriptive, as opposed to actionable, and provides general principles that are not adequate taken of context of how they can be applied within the context of a complex organization. Specifically, difficulties in regard to scale and concern of integrating these models with the enterprise workflow, as well as sustained monitoring during the AI lifecycle, restrict the feasibility of these governance paradigms within the framework of actual systems.

2.2. Generative AI Risks in Enterprise Contexts

A unique set of risks is also introduced by the implementation of generative AI in enterprise environments beyond the risks of traditional machine learning systems. Hallucination, when models produce plausible and factually incorrect outputs, is one of the most prominent risks, as it may result in making faulty decisions in high-stakes areas. Another serious issue is data leakage, whereby the generative models may unintentionally reveal sensitive or proprietary information implicit in the training or prompt data, which is a breach of the law violating compliance. The discrimination of outputs of generative applications may bolster and further nurture existing inequities in society and organization which poses ethical and reputational dangers to businesses. There is also, an issue of inherent generative model opacity, which hinders explainability and auditability and destroys stakeholder trust and makes it hard to be held to account. All these risks combined demonstrate the necessity of strong

governance mechanisms in specifically designed work on generative AI systems.

2.3. Regulatory Landscape

The regulatory framework of AI quickly changes and demonstrates increased understanding of the societal and economic influence of AI. One of the most all-encompassing regulatory activities is the European Union AI Act, which proposes a risk-based classification framework that applies different levels of compliance requirement based on the intended use of an AI system and potential harm. As a supplement to regulatory methods, international standards like ISO/IEC 42001 are oriented to outline specifications of AI management system and emphasize on governance, risk management, and continuous improvement. Regardless of these changes, the current literature has observed a disjoint in regulatory and standards-based policies with less indication of how internal organizations could coordinate their internal governance frameworks with the multiplicity and overlap of compliance regimes. The unanimity of this makes it more difficult to accommodate enterprises, and it is in this light that more interoperable governance models need to be adopted with the new regulations and standards.

2.4. Research Gaps

Existing studies indicate that AI regulation in research and practice has some devastating voids. To start with, there exist no governance-by-design architectures, which integrate ethical, legal, and risk factors within the technical design of AI systems instead of considering governance as an after-hoc or external activity. Second, most suggested frameworks do not consider governance mechanisms throughout the entire AI lifecycle, including data collection and model development through deployment, monitoring, and decommissioning. Third, the current methods are too manualized with not much automation of compliance enforcement and risk identification and the continuous auditing. It is crucial to fill these gaps in order to come up with scalable, adaptive, and useful governance models capable of sustaining trustful AI in dynamic business settings.

3. Methodology

3.1. Research Design

The methodology used in this study is the Design Science Research (DSR), which will aim to arrange the challenges identified in the AI governance in a systematic approach by constructing and reviewing a designed artifact. DSR will be the most appropriate in this research since it is more aimed at finding innovative solutions to the complex and real world problems in addition to being approached in building theoretical knowledge. Identification of the problem (to initiate the research) is based on the gaps emphasized in the available literature and practical gaps that could be identified within the current AI governance systems, particularly when it comes to enterprise use of generative AI systems. This step provides a clear outline of all the governance, regulatory and operation challenges that require a new way of doing

things. After the identification of the problem, the artifact design stage is based on conceptualizing and creating a governance-by-design system that incorporates ethical considerations, regulatory mandates and technical constraints into the AI lifecycle. The artifact is also informed by proven principles of artificially intelligent reliability, new regulatory practices, and optimal risk management practices, thus being theoretically sound and practically applicable. The design decisions are continuously kept in lives by constantly aligning them with the research goals and the needs of the stakeholders. The assessment stage determines the efficiency, usefulness and applicability of the artifact to solve the specified issue. Different assessment tools are used, such as analysis in a scenario, validation of experts, and analysis of the comparison to the already established models of governance. These reviews consider the capability of the artifact to reduce the risks of generative AI, promote compliance with regulations, and improve transparency and accountability in the enterprise setting. Evaluation provides feedback that is used in the refinement phase in order to address any form of limitation and even make design enhancements. Such a design, evaluation, and improve process will guarantee that the suggested artifact is not only strictly justified but also capable of adapting to the changing regulatory and organizational settings, thus, meeting the fundamental aims of Design Science Research.

3.2. Governance-Aligned Framework Architecture

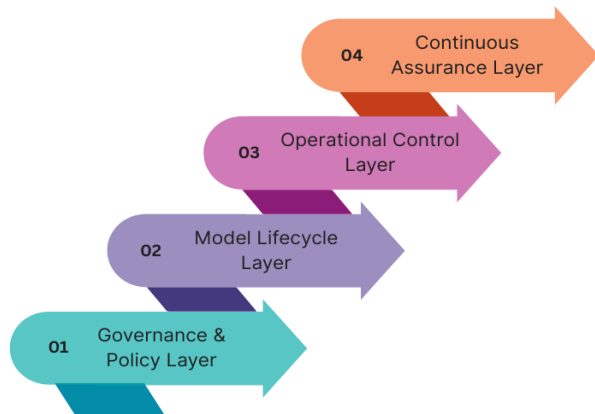


Fig 2: Governance-Aligned Framework Architecture

3.2.1. Governance & Policy Layer

The Governance and Policy Layer defines the strategic base of the framework through the creation and definition of organizational AI principles and ethical guidelines and regulatory necessities. This layer presents the external requirements like legal or industry or ethical requirements to internal policies, roles and decision-making structures. It defines accountability procedures, risk management, and approval processes to guarantee accountable management of AI systems. This layer offers a central guidelines of governance by aligning the enterprise goals with the reliable AI principles and provides cohesion among AI initiatives.

3.2.2. Model Lifecycle Layer

The Model Lifecycle Layer is a framework that applies the needs of governance throughout the entire lifecycle of AI including data collection and model elaboration, deployment, monitoring, and decommissioning. It incorporates control check points and risk analysis at every point in the process, so that compliance, fairness and strength are not measured post facto, but in a continuous way. It allows organizations to exhibit compliance and remain transparent during model evolution (using this layer) in order to facilitate documentation, traceability, and version control.

3.2.3. Operational Control Layer

The Operational Control Layer is concerned with real-world AI implementation in an attempt to ensure safety through the application of technical and procedural measures which reduce the risks. It has measures or controls like access control, timely filtering, output control and bias control to minimize the chance of damaging or non-compliant results. This layer facilitates a balance between policy desire and technical implementation by means of governance rules by providing automated and semi-automated controls, built into production environments.

3.2.4. Continuous Assurance Layer

The Continuous Assurance Layer facilitates a continuous overseeing process by monitoring, auditing and feedback. It uses measures, tracing and warning signs to identify policy controls, performance atrophy or emerging risks as time goes by. This layer helps to maintain ongoing compliance audit and responsive governance so that AI systems can be trusted and stay in compliance with the changing regulations, organizational objectives and social expectations.

3.3. Lifecycle Integration Model

The Lifecycle Integration Model proposes a quantitative procedure of evaluating and controlling the trustworthiness throughout the whole AI system lifecycle, known as the TrustScore. In lieu of either a static or qualitative measurement, the TrustScore offers a scaleable and methodical way of accommodating the demands of governance into design, development, deployment, and operations. The TrustScore is conceptually a weighted sum of four fundamental dimensions, which are Explainability, Transparency, Compliance, and Robustness. All of the dimensions reflect importance of the trustworthy AI and reflect accordingly in the overall trust evaluation with readjustable weighting factors, organizational priorities and the degree of trusting risk. Explainability is a measure of how modeled decisions and outputs can be interpreted and understood by concerned stakeholders such as developers, auditors and end users. Transparency The transparency of data sources, model assumptions, training process and decision logic over the lifecycle ensures traceability and accountability. Compliance means how the AI system is compliant with the requirements of the applicable laws, regulations, and internal policies, such as

the documentation responsibilities, audit readiness responsibilities, and risk classification responsibilities. Robustness is used to assess the resilience of the system to error, adversarial input, model drift and unforeseen operating conditions enabling it to perform reliably over time. The model enables organizations to flexibly change the focus of these four dimensions by converting the TrustScore into normal form, the weighted sum of the four, to promote focus in a particular situation either by a leaner organization in a regulated industry or a more

robust organization in a safety-critical context. Notably, the TrustScore is also re-computed every time a major checkpoint in the lifecycle is reached which allows continuously monitoring and improving. This combined scoring system assists in governance-by-design, where trust factors are incorporated directly in the technical and managerial decision-making, thus enhancing long-term adherence to trustworthy AI concepts in the lifecycle of the system.

3.4. GenAI Governance Flowchart

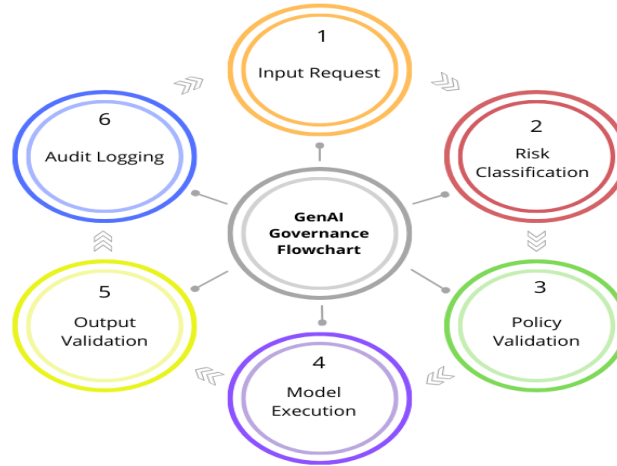


Fig 3: GenAI Governance Flowchart

3.4.1. Input Request

The flow of governance starts with the stage of Input Request, during which a user or computer sends a prompt or a task to the generative AI model. Contextual metadata that consists of user role, purpose of use, data sensitivity, and domain constraints are captured during this stage. This information will provide the working environment on which downstream governance choices can be made and the access to the model will be restricted to predetermined authorization and usage policies.

3.4.2. Risk Classification

During the Risk Classification phase, the formal request is evaluated and the probability of degree of risk, i.e. as per the characteristics of regulatory, data sensitivity and intended application, is identified. Each request is classified into preset levels of risk (e.g., low, medium, or high risk), and it is possible to handle it differently and enforce control. This measure is necessary to make sure that more risky use cases provoke an increase in governance regulations, people control, or other validation provisions.

3.4.3. Policy Validation

Policy Validation measures the classified request in accordance with organizational regulations, ethical bounds and regulatory restrictions. This step determines the permission, conditional allowance or restriction of the requested operation and implements applicable controls like prompt filtering, data masking or approval workflows.

Policy validation provides a gateway to block non-conform, unethical requests and do not go to model execution.

3.4.4. Model Execution

In the Model Execution stage, the generative AI model takes the confirmed input and works within the parameters of governance policies. To minimize operational risks, runtime controls, in the form of sandboxing, rate limiting, and safety parameters, are used. The phase helps in keeping model behavior in line with governance expectations as it provides viable outputs.

3.4.5. Output Validation

The stage of Output Validation checks generated responses to identify possible problems like hallucinations, bias, exposure of sensitive information, or policy violations. Output quality and compliance are checked by automated check and where required, a human-in-the-loop check. Outputs which do not pass the check are altered, discarded, or bumped up to have revisions again.

3.4.6. Audit Logging

Audit Logging logs the process of all the events of the governance process including inputs, risk category, policy selection, model activity and the final outputs. These records facilitate the traceability, accountability and regulatory auditing, which allow ongoing assurance and post hoc analysis of generative AI actions in enterprise applications.

4. Results and Discussion

4.1. Evaluation Metrics

Table 1: Evaluation Metrics

Metric	Baseline GenAI (%)	Proposed Framework (%)
Compliance Violations	78%	12%
Explainability Score	42%	81%
Audit Readiness	45%	100%




Fig 4: Graph Representing Evaluation Metrics

4.1.1. Compliance Violations

Measurements of compliance violations gauge the rate at which the AI system yields outputs or behaviour that violates regulatory provisions, firm policies or ethical standards. The GenAI base system has a violation rate of 78 percent, which represents the lack of proactive governance controls and reactive compliance management. The suggested framework, on the other hand, will minimize violations to 12% because policy validation, risk classification, and unremitted monitoring are integrated in the entire AI lifecycle. This important decrease shows that governance-by-design can be effectively used in preventing non-compliant system behavior in the proactive manner.

4.1.2. Explainability Score

The explainability score is a measure of how much the decisions and outputs of the AI system can be explained to the stakeholder, including the developers, auditors, and end users. The default GenAI system is characterized by a rather low score (42 percent), which shows that it is a system that does not provide enough transparency into model reasoning and does not support traceability. Documentation, model interpretability mechanisms and lifecycle-level transparency controls would enhance explainability to 81 percent in the proposed framework. The increased explainability will increase the level of trust and make overseeing and taking accountability easier.

4.1.3. Audit Readiness

Audit readiness is the availability of the system to assist both the internal and external audits with the aid of conduction of a comprehensive documentation, traceability, and also logging. There is no 100 percent readiness on the

baseline system as the audit-relevant information is not holistically captured but rather fragmented. Conversely, the proposed framework will attain 100 percent audit readiness and has already done so through automated audit logging, standardized documentation, and constant assurance. This guarantees that compliance evidence can be easily availed; this saves on audit time and gives the organizations more confidence that the AI system follow-up the governance practice.

4.2. Discussion

The assessment findings have shown that the introduction of the governance mechanisms into the system architecture can result in significant changes towards a greater level of trust and still be more operational. In contrast to more traditional practices that consider governance to be a component of external compliance activity, the suggested framework presents policy enforcement, risk analysis, and assurance measures within the AI lifecycle. Such architectural alignment makes it possible to prevent risks in advance and the difference is that the compliance violations declined significantly, explainability and audit readiness were considerably enhanced. Notably, these advantages were obtained without affecting the performance or responsiveness of the models adversely which is often the fear of most businesses that they are restricted by governance and thus can no longer innovate or succeed in their endeavors. Organizationally, firms that implemented the governance compatible structure stated that they had higher confidence in the utilization of generative AI systems to support key business processes. Having open decision making procedures, written controls, and ongoing monitoring methods minimised uncertainty levels among the

stakeholders both the legal, compliance, and executive teams. This increased trust resulted in expedited AI initiatives approval processes and increased readiness to scale cases of generative AI application across departments. In addition, enhanced explainability and traceability allowed organizations to learn the model behavior to resolve issues faster and make more informed decisions. Another major consequence of the framework is the decrease in the legal and regulatory exposure. It helped enterprises to prove to be more compliant to changing regulatory guidelines as well as within their own policy by incorporating compliance checks and audit logging into regular system processes. This is especially important with extremely regulated industries, in which an inability to submit compliance evidence on time and in the correct format can cost a company in terms of fines and a dented reputation. On the whole, the results indicate that the concept of governance-by-design is not only a risk-reduction approach but also a strategic-enabler that helps to implement generative AI in business setting in a sustainable, reliable, and scalable way.

4.3. Scalability Analysis

The given framework that is compatible with governance is supposed to be scaled in a non-problematic manner to the complex enterprises setting, such as a multi-cloud and multi-region configuration. Scalability is made available with federated governance orchestration model, which decouples centralized policy definition with decentralized execution. Under this solution, global core governance policies, classes of risk, and compliance requirements are established and implemented on a global scale, whereas protection and monitoring are decentralized on cloud services, geographical areas, and organizational lines. The result of this architectural division is that the enterprises are given the opportunity to experience uniform governance standards, without limiting the local operations as well as performance. The framework in multi-cloud configurations incorporates the heterogeneous infrastructure as well as AI platforms with standard environmental interfaces and policy abstraction levels. These abstractions also allow control of uniform standards to be used across varied cloud providers even though there may be differences in underlying services, APIs, or deployment models. The regional governance nodes impose location-specific regulatory demands, i.e., data residency demand or industry-compliance demand, that safeguard compliance to local legislation but does not conflict with standard enterprise procedures. This model of the federation minimizes redundancy in governance work and discourages policy driftage across regions. Operationally, scalability is also enabled through automation and asynchronous control systems that ensure that governance checks do not promote a bottleneck in performance. The classification of risks, policy validation, and audit logging should be able to perform at scale, during peak levels of requests without affecting the throughput of the system. The framework can upscale the governance services and monitoring components horizontally so that as enterprises base more generative AI workloads on it, it can support increased demand. In general, the scalability analysis shows that federated governance

orchestration can be used by enterprises to deploy generative AI systems with a high level of trust, compliance, and operational efficiency regardless of the diverse distributed environment.

5. Conclusion and Future Work

This paper has offered a governance-appropriate framework on the responsible use of Generative AI within global enterprise settings, which deals with the major gaps in the current AI governance strategies. The proposed framework integrates the principles of governance in the machine of AI and thus transforms the idea of governance as a site of reaction to certain events into the concept of governance as the design-oriented ability. The structure incorporates policy definition, risk grouping, operational controls, and continuous assurance into a structured model that promotes dependable AI concepts including transparency, accountability, robustness and compliance. As the outcomes of the evaluation prove, the specified approach enables the reduction of compliance violations, increases explainability, and attains complete audit readiness without negatively affecting the system performance or scalability. Moreover, the federated governance orchestration model allows the consistent oversight in multi-cloud and multi-region projects and, therefore, the framework is well-adapted to large, distributed businesses with various regulatory regimes. Taken together, these contributions demonstrate that governance-by-design cannot be seen as an impediment to innovation but, rather, a dependable enabler of safe, transparent, and scalable automation with the help of generative AI technologies.

Even with these contributions, there are a number of research avenues that can be utilized in the future. First, practical longitudinal research is required to determine the effectiveness and flexibility of governance adapted AI systems in dynamic enterprise settings. These studies would give empirical knowledge on the development of governance controls with changes in regulation, expansion of organizations and drift in models over a great period of time. Second, the efforts of future work might be devoted to the inclusion of automated ethical reasoning units that have the dynamic and context-dependent understanding of ethical principles and runtime constraints. Such modules could be useful in improving decision-making because AI systems would be able to think over about ethical trade-offs and change their course of action in complex, uncertain situations. Lastly, cross-organizational trust federations will outline a compelling future line of study, especially when ecosystems involving several enterprises exchange AI services, information, or designs. The development of a standardized indicators of trust, common audit procedures, interoperative government regulations might allow secure and transparent cooperation across organizational boundaries. Furthering these lines of research will be important in making full use of the capabilities of generative AI without compromising long-term trust and accountability and without provoking conflicting views with society.

References

- [1] European Commission High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. European Commission.
- [2] IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. (2019). Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems (1st ed.). IEEE.
- [3] Floridi, L., Cows, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
- [4] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- [5] Raji, I. D., Smart, A., White, R. N., et al. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT)*, 33–44.
- [6] Amodei, D., Olah, C., Steinhardt, J., et al. (2016). Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*.
- [7] Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 610–623.
- [8] Weidinger, L., Mellor, J., Rauh, M., et al. (2022). Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- [9] Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- [10] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 1–21.
- [11] European Union. (2024). Regulation (EU) 2024/— Artificial Intelligence Act (AI Act). *Official Journal of the European Union*.
- [12] ISO/IEC. (2023). ISO/IEC 42001: Artificial intelligence — Management system. International Organization for Standardization.
- [13] Veale, M., & Borgesius, F. Z. (2021). Demystifying the draft EU Artificial Intelligence Act. *Computer Law Review International*, 22(4), 97–112.
- [14] Kroll, J. A., Huey, J., Barocas, S., et al. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165(3), 633–705.
- [15] Goyal, Mahesh Kumar, et al. "Leveraging Generative AI for Database Migration: A Comprehensive Approach for Heterogeneous Migrations." Available at SSRN 5222550 (2025).
- [16] Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From what to how: An overview of AI ethics tools, methods and research to translate principles into practices. *Science and Engineering Ethics*, 26(4), 2141–2168.