ISSN :2584-2706

AI Governance, Ethics, and Responsible Innovation

Zouhair Chiban
Phd Student in Private Law
Faculté des Sciences Juridiques, Économiques et Sociales
Université Abdelmalek Essaâdi, Tanger, Maroc

Abstract
In spite of the fact that there is widespread agreement regarding the ethical principles that need to guide artificial intelligence, there is still a significant gap between these normative objectives and the operational realities of technical implementation. An investigation into how to operationalize a proactive, iterative AI governance paradigm is the focus of this study. The investigation is conducted via the prism of private law liabilities and constant technical verification. Both corporate compliance theater and opportunistic ethics shopping are the driving forces behind our inquiry. Both of these shortcomings in practice have occurred repeatedly.

During the course of the study, a multi-layered methodological framework is utilized. This framework triangulates a comparative assessment of growing international regulatory frameworks and a "cross-sectoral typological mapping of auditing procedures from banking, data privacy, and environmental law." The findings reveal a systemic disconnect: corporate ethical commitments continue to be largely detached from operational audit systems, leaving data scientists frustrated by guidelines that are too abstract to act on. Furthermore, existing environmental, social, and governance standards do not provide any AI-specific tools for quantifying the hidden socio-technical harms that are embedded in global supply chains. rules that are abstract, whereas the existing ESG standards do not provide the AI-specific guidance that is required to assess the hidden socio-technical damages that are present throughout global supply chains.The purpose of this article is to offer a solution to this problem by advocating for the transition from voluntary soft law to an enforceable, longitudinal regulatory framework. We propose that corporate governance frameworks be merged with automated auditing pipelines, value-based engineering (also known as ethics-by-design), and independent third-party oversight. This would result in the transformation of abstract ethical concepts into private law duties that can be verified.

Keywords: *AI Governance, Ethics-by-Design, Algorithmic Auditing, Responsible Innovation, ESG Metrics, Corporate Liability, Private Law.*

I. INTRODUCTION

The rapid proliferation of intelligent systems demands robust guardrails to ensure that deployment respects foundational human norms and emerging regulatory mandates. At the core of this aim is the deployment of **multi-stakeholder**[1] **collaborative models** [2]that combine **ethical principles** [3]with **robust oversight mechanisms** [4]across the entire development life cycle (Batool 2025).

A holistic approach of this nature requires the use of **systematic frameworks** [5]which embed regulatory compliance and ethical accountability at the beginning stages of a project, beyond this proactive integration it is also imperative to move from singular algorithmic adjustments to a focus on **comprehensive system-level patterns**[6] for operationalizing these standards (Lu et al 2023). Rather than treating governance as a one-time compliance exercise, organizations that approach it as an iterative process — supplementing existing frameworks with responsibly collected real-time data — are better positioned to monitor system behavior

[1] **Multi-stakeholder:** A governance approach that actively involves a diverse range of actors including policymakers, legal experts, tech developers, and civil society in the decision-making process to ensure balanced regulation.
[2] **Collaborative models:** Structural frameworks designed to facilitate shared responsibility, data exchange, and joint oversight between different institutions and sectors to manage complex autonomous systems.
[3] **Ethical principles:** Clear moral guidelines like fairness, honesty, and respecting human rights—that software developers must follow to make sure their technologies do not cause harm or discriminate against users.
[4] **Robust oversight mechanisms:** Strong and reliable control tools such as constant technical checks and independent human reviews used to monitor a system continuously and stop it if it makes an error or breaks a law.
[5] **Systematic frameworks:** Step-by-step organized plans or blueprints that companies follow from start to finish to ensure they obey laws and respect ethical rules at every stage of a project.
[6] **Comprehensive system-level patterns:** Looking at the behavior of a whole technology network as a single big picture, instead of just checking small, isolated pieces of code or single algorithms.

ISSN :2584-2706

and address risks like systemic bias and misinformation before they escalate. This paradigm also calls on policymakers and industry leaders to embed norms of transparency and accountability as a way of navigating the socio-technical tensions inherent to generative technologies (Ahmad et al. (2024) This shift toward **ethics-by-design**[7] — characterized by **human-centric oversight** [8]and less of the **reactive regulatory measures** [9]allows development trajectories to remain responsive to evolving threats.

By employing computational governance approaches such as **automated auditing** [10]and **real-time monitoring**[11]**,** organizations can transition from isolated silos to a more interconnected, global and scalable system. This requires an **ecosystem-level approach** [12]: multi-stakeholder collaboration supported by a broad transparency architecture — including **machine-readable ethics labels** [13]and **algorithmic impact assessments** [14]— to bridge the gap between high-level principles and practical technical documentation.

Practitioners will need to translate the changes proposed in **normative ethical standards** [15]into concrete action, with particular focus on tools such as ethical impact assessments and **bias auditing**[16], On the one hand, these technical strategies are backed up by independent **third-party auditing processes** [17]that are meant to objectively certify operational performance in relation to externally set **safety benchmarks** [18](Cheong I. et al, 2024).Modular governance frameworks further enable organizations to take broad ethical guidelines and tailor them to the specific demands of their sector.

Recent literature shows that there is a complete gap between broad agreement on core ethical principles: justice, transparency and accountability and getting it to practical implementation (Irhas, 2026; Agarwal & Nene, 2025). This gap is compounded by compliance theatre — performative adherence to organizational verification and safety cultures that conceals a lack of sustained, substantive accountability (Shneiderman, 2020). Scholars suggest that enforceable governance practices and independent audit frameworks are necessary to move accountability from abstract, principle-based policies toward concrete **technical and organizational verification**[19], which hopefully mitigates this problem (Falco et al., 2021).

Mökander and Floridi (2022) observe that organizations adopting **ethics-based auditing**[20] as a continuous, proactive practice are better equipped to move from theoretical alignment to verifiable risk governance; integrating these capabilities into existing workflows transforms auditing from a box-ticking formality into a genuinely useful operational tool (Mökander & Floridi, 2022). This also underscores the need for bottom-up case studies and localized practices that can complement top-down regulatory frameworks — ensuring implementation remains fit for purpose across diverse technological and institutional contexts (Mittelstadt, 2019). This **methodological shift**[21] calls for combining technical tools — such as open-source bias mitigation libraries — with non-

[7] **Ethics-by-design:** Building legal rules and moral values directly into the software code from the very first day of development, so the technology automatically respects human rights without needing to be fixed later.
[8] **Human-centric oversight:** Keeping a human expert in control of the technology, giving them the final power to review, correct, or completely stop an AI's decision if it harms people or violates the law.
[9] **Reactive regulatory measures:** Laws or penalties that governments pass only *after* a technology has already caused a major scandal, accident, or harm to the public, rather than preventing the problem in advance.
[10] **Automated auditing:** Using specialized computer programs to continuously inspect and test an AI's code, logs, and decisions automatically to make sure it is following legal standards without requiring a human to check everything by hand.
[11] **Real-time monitoring:** The continuous, live watching of an AI system while it is running, allowing software or human supervisors to instantly see failures, errors, or unexpected behaviors the exact second they happen.
[12] **Ecosystem-level approach:** Looking at technology not as a single tool, but as a large connected network where tech companies, users, laws, and social impacts all interact together and depend on one another.
[13] **Machine-readable ethics labels:** Digital tags or certificates embedded inside a software's code that allow other computers to instantly read, understand, and verify the safety rules and ethical limits of that technology.
[14] **Algorithmic impact assessments:** Official review reports that companies must complete before launching an AI, designed to check for potential risks such as discrimination, privacy leaks, or negative social consequences.
[15] **Normative ethical standards:** The official, agreed-upon rules and core values that define how an AI system *should* ideally behave to be fair, honest, and just, setting the baseline for what is right and wrong.
[16] **Bias auditing:** A technical review process that checks an AI's code and data to see if it makes unfair decisions or discriminates against certain groups of people based on race, gender, or background.
[17] **Third-party auditing processes:** Independent inspections carried out by external, neutral experts to verify that a company's technology complies with safety, legal, and ethical standards without any conflict of interest.
[18] **Safety benchmarks:** Standardized, industry-wide tests and evaluation metrics used to measure an AI system's security, stability, and resistance to errors before it can be legally approved for public use.
[19] **Technical organisational verification:** A dual-layer control process that checks both the software code (technical) and the company's internal management structures (organisational) to guarantee full compliance with safety laws.
[20] **Ethics-based auditing :** A continuous evaluation process that measures whether a technology's development, data usage, and decisions respect human rights, fairness, and moral values.
[21] **Methodological shift :** A fundamental change in how organizations approach a problem, moving from purely theoretical policies to a practical combination of technical and non-technical tools.

ISSN :2584-2706

technical instruments, including codes of practice and contractual conditions, to build a robust **evidentiary trail**[22] (Eitel-Porter, 2020),Technology providers must also conduct governance audits that examine whether internal decision-making structures, incentive systems, and management practices give adequate executive backing to risk-mitigation strategies (Mökander et al., 2024). This shift has highlighted an urgent need for comparative research into how ethical principles are interpreted and translated into action in different organizational and geographical settings (Ahmed, 2025). This calls for **adaptive risk governance frameworks** [23] that can keep pace with accelerating technology (Cheong B. C., 2024). Ultimately, transforming auditing into a comprehensive process that accounts for the long-term effects of systems on users, society, and the environment (Mökander, 2023) is central to any credible path forward for AI governance.

II. METHODS

This study uses a "**multi-layered analytical**[24] framework, which is intended to measure the effectiveness of existing AI governance mechanisms by transforming qualitative ethical narratives into quantifiable regulatory and technology standards (Mökander et al., 2024). Given the **socio-technical nature**[25] complexity of algorithmic ecosystems, the methodology deliberately moves beyond static, self-reported compliance documentation. It draws instead on a "**dual-track analytical approach,**[26]" namely a comparative legal review of novel transnational AI regulatory mandates and a **cross-sectoral typological mapping**[27] of effective compliance mechanisms (Birhane et al. 2024).

Initially, an investigative document review of corporate governance protocols (a review undertaken to analyze algorithmic-use risk analyses and institutional policy statements on risk mitigation prior to deployment; Solanki et al., 2022) was performed. This textual data is then assessed using a risk management kernel framework that maps corporate policies across five pillars: risk identification, collection, assessment, compliance, and monitoring, each adapted to the specific challenges of the AI lifecycle.

Second, in response to the limited technical operationalization available in current frameworks, this research classifies existing AI-specific monitoring instruments into a comparative typology. This typology juxtaposes innovative algorithmic auditing methods with traditional regulatory approaches that have existed for decades in three well-established areas: environmental science, data privacy, and financial law (Ayling & Chapman, 2021). Through this cross-sector parallel analysis, we identify structural legal mechanisms — **standardized professional accreditation**[28]", continuous independent oversight, and **life-cycle impact reporting**[29] that we can repurpose to hold AI development accountable through this "cross-sector parallel analysis".

The framework also examines how technical audits can be integrated into corporate **Environmental, Social and Governance (ESG) reporting structures**[30]. Differentiating between technology-defined impact assessment and internally determined organizational governance audits, the framework explores how investors and regulators can utilize ESG metrics as a long-term corporate responsibility compliance measure (Minkkinen, Niukkanen et al. 2022). Finally, the framework evaluates the structural viability of continuous real-time monitoring — particularly its advantages over traditional **point-in-time snapshot audits** [31]and its resistance to internal conflicts of interest in corporate recordkeeping (see Minkkinen, Laine et al., 2022; Raji et al., 2022).

III. RESULTS

[22] **Evidentiary trail :** A secure, chronological record of documents, test results, and decisions that proves a company followed all legal and ethical steps during a technology's creation.
[23] **Adaptive risk governance frameworks :** Flexible legal and management structures designed to automatically update, learn, and evolve in response to rapid technological advancements.
[24] **Multi-layered analytical framework:** A research methodology that examines a complex problem from multiple distinct angles simultaneously, such as combining legal, technical, and organizational perspectives.
[25] **Socio-technical nature:** The complex reality where social human behavior, legal rules, and computer technologies heavily influence and shape one another, making it impossible to separate the code from its social context.
[26] **Dual-track analytical approach:** A research strategy that runs two separate investigations at the same time—such as a legal analysis and a technical mapping—to arrive at a balanced, comprehensive conclusion.
[27] **Cross-sectoral typological mapping:** The process of identifying, classifying, and comparing different types of compliance tools across various business industries (like banking, health, and automotive) to find universal solutions.
[28] **Standardized professional accreditation:** An official licensing system where auditors must pass state-approved exams and get certified before they are legally allowed to inspect AI systems.
[29] **Life-cycle impact reporting:** The legal obligation for a company to publish regular, transparent reports about a technology's risks at every single stage, from its initial coding to its final deletion.
[30] **Environmental, Social, and Governance (ESG) reporting structures:** Official frameworks used by companies and investors to evaluate a business's sustainability, ethical footprint, and social responsibility alongside its financial profits.
[31] **Point-in-time snapshot audits:** Traditional inspections that only check a company's systems on one specific day of the year, failing to see what happens before or after that exact date.

ISSN :2584-2706

The analysis reveals that while organizations increasingly adopt high-level ethical principles, these are rarely embedded in structural operational audit systems — leaving a persistent gap between corporate intent and actual development practice (Bartsch et al., 2025). The prevailing tendency on a policy level to rely on **soft-policy approaches**[32] significantly worsens this discrepancy. Voluntary frameworks offer no precise, legally binding standards against which technical compliance can be measured, leaving companies uncertain about how to reconcile competing human values and commercial pressures at the very stages of the product lifecycle where implementation decisions are made. The data also illustrates **ethics shopping**:[33] firms selectively endorse voluntary guidelines that accommodate existing practices, revealing the systemic risks of leaving governance frameworks entirely unregulated (Mökander, Axente et al., 2021).

A significant structural weakness found in the cross-sectoral research concerns business sustainability reporting. As of October 2023, both identified Environmental, Social and Governance (ESG) frameworks offer generic guidance, but none are specific to AI. As such, there are no standards for reporting algorithmic indicators that would help quantify these indirect, concealed harms called global **digital value chains**[34] — from the massive energy consumption of data processing to the human exploitation involved in dataset labelling.

The findings point to a strong imperative from both technical practitioners and legal scholars for enforceable, sector-specific regulations that translate abstract principles into concrete, context-relevant legal norms to counter these systemic shortcomings (Morley et al., 2021).

Data scientists expressed significant frustration with high-level ethical codes, finding that abstract principles rarely translate into the mathematical and technical realities of day-to-day coding (Morley et al., 2019). On top of that, there is also a significant lack of objective, automated assessment tools available to companies.

The findings suggest that resolving this ambiguity requires a shift toward **value-based engineering**[35] — moving beyond conventional architectural assumptions to embed ethical protections directly into the technical design process. Participatory design methods can also help resolve these tensions more durably, by bringing software developers, compliance officers, and external stakeholders to agreement on what counts as verifiable.

IV. DISCUSSION

This study argues that AI governance discourse must move away from treating ethics as an upfront hurdle to clear before deployment, and instead treat it as something woven continuously through the entire system lifecycle — from development through deployment and beyond — mirroring the way liability flows uninterrupted across globalized supply chains. Avoiding **ethics washing**[36] where voluntary frameworks are used by companies to forestall binding regulation while leaving problematic practices intact (Ibitoye et al., 2025) — requires more than rhetorical commitment.

Organizations must develop a granular taxonomy of technical tools and standardized metrics. Current guidelines are, we argue, too abstract to meaningfully guide data scientists navigating competing pressures; the vagueness of many high-level principles routinely obscures socio-political harms and structural inequities behind ostensibly neutral language (Alahmadi, 2025; Hickok, 2020).

From a private law perspective, closing this operational gap requires more than technical fixes — it demands a structural rethinking of corporate governance and liability distribution. Regulatory regimes must specify clear legal obligations calibrated to the actual degree of control an entity holds over a model's outputs, data inputs, and commercial deployment.

Global companies must move beyond siloed management structures and establish independent, multidisciplinary oversight boards capable of translating broad normative commitments into concrete technical architecture — bridging the divide between abstract ethical principles and long-term legal accountability (Munn, 2022; Minkkinen et al., 2022). For these internal governance bodies to thrive within an organization they must be structurally supported by an automated assessment pipeline that regularly and objectively deliver feedback to developers — directly linking legal compliance requirements to system architecture. Drawing on established regulatory models from environmental and financial law — mandated Algorithmic Impact Assessments and

[32] **Soft-policy approaches:** Regulatory methods that rely on voluntary guidelines, codes of conduct, and non-binding recommendations instead of strict, legally enforceable penal laws.
[33] **Ethics shopping:** A deceptive corporate practice where a company selects and follows only the easiest, non-binding ethical guidelines that match its current behavior, while ignoring stricter, more important rules.
[34] **Global digital value chains:** The international network of labor, data processing, and hardware infrastructure required to build, train, and maintain an AI system across different countries.
[35] **Value-based engineering:** An engineering methodology where human values, ethical principles, and legal rights are treated as strict technical requirements during the system design phase, just like speed or memory capacity.
[36] **Ethics washing:** The practice of using superficial ethical campaigns, statements, or voluntary codes to create a positive public image, aimed at distracting from a company's actual lack of legal compliance or harmful practices.

ongoing third-party audits — would help anchor accountability in codified, evidence-based practice (Birkstedt et al., 2023). Rather than addressing software products in isolation, private law could draw on established risk management instruments — such as human rights impact assessments — to build procedural legal norms that treat AI systems as the complex, high-stakes business assets they are. Such institutional redesign would elevate ethics-by-design from a regulatory slogan to a verifiable standard of corporate due diligence — one where systemic failures give rise to tortious liability across the full AI development lifecycle (Jobin et al., 2019).

V. CONCLUSION

Ultimately, effectively addressing the failures of AI ethics requires moving beyond non-binding, voluntary codes toward enforceable regulatory environments with built-in mechanisms of legal accountability (Akinrinola et al. 2024; Hagendorff 2020).

This shift relocates responsibility from isolated individuals to collective structures embedded within business models (Häußermann & Lütge, 2021) and in doing so, demands a deeper transformation of organizational culture. It can also help resolve the longstanding ambiguity that has rendered principles-based ethics codes largely ineffective in practice, by recasting normative objectives as verifiable engineering requirements (Griffin et al., 2023; Krijger et al., 2022).

By pairing these governance structures with iterative, ethics-based auditing, organizations can continuously assess conformance to core principles and address gaps throughout the development lifecycle (Amugongo et al., 2023). This means that meaningful escalation mechanisms are essential: research of this kind should drive tangible product changes, not gather dust as static reports filed away and forgotten.

Beyond internal controls, policymakers will need to incentivize the adoption of elevated corporate governance frameworks — ones that mandate independent audits for high-risk applications — thereby filling the institutional gap between voluntary guidelines and enforceable conformity assessment (Ball & Helliwell 2025).

REFERENCES

[1]. Agarwal, A. K., & Nene, M. J. (2025). A five-layer framework for AI governance: integrating regulation, standards, and certification. *Apress eBooks*, 19(3), 535–555.
[2]. Ahmad, N., Anshari, M., Hamdan, M., & Ali, E. (2024). Ethics and public trust in AI governance: A literature review. *International Journal of Law Government and Communication*, 9(37), 296–305.
[3]. Ahmed, I. (2025). Navigating ethics and risk in artificial intelligence applications within information technology: A systematic review. *American Journal of Advanced Technology and Engineering Solutions*, 1(1), 579–601.
[4]. Akinrinola, O., Okoye, C. C., Ofodile, O. C., & Ugochukwu, C. E. (2024). Navigating and reviewing ethical dilemmas in AI development: Strategies for transparency, fairness, and accountability. *GSC Advanced Research and Reviews*, 18(3), 50–58.
[5]. Alahmadi, D. (2025). AI Ethics Principles for AI Developers in Academia: A Critical Assessment and Analysis. *TEM Journal*, 2514–2529.
[6]. Amugongo, L. M., Kriebitz, A., Boch, A., & Lütge, C. (2023). Operationalising AI ethics through the agile software development lifecycle: a case study of AI-enabled mobile health applications. *AI and Ethics*, 5(1), 227–244.
[7]. Ayling, J., & Chapman, A. (2021). Putting AI ethics to work: are the tools fit for purpose? *AI and Ethics*, 2(3), 405–429.
[8]. Ball, B., & Helliwell, A. C. (2025). The 3D model of ethical AI practice. *AI and Ethics*, 5(6), 5955–5969.
[9]. Bartsch, S. C., Nguyen, L. H., Schmidt, J.-H., Du, G., Adam, M., Benlian, A., & Sunyaev, A. (2025). The Present and Future of Accountability for AI Systems: A Bibliometric Analysis. *Information Systems Frontiers*, 27(6), 2463–2484.
[10]. Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: a systematic literature review. *AI and Ethics*, 5(3), 3265–3279.
[11]. Birhane, A., Steed, R., Ojewale, V., Vecchione, B., & Raji, I. D. (2024). AI auditing: The Broken Bus on the Road to AI Accountability. *IEEE satml*, 612–643.
[12]. Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: themes, knowledge gaps and future agendas. *Internet Research*, 33(7), 133–167.
[13]. Blanchard, A., Thomas, C. E., & Taddeo, M. (2024). Ethical governance of artificial intelligence for defence: normative tradeoffs for principle to practice guidance. *AI & Society*, 40(1), 185–198.
[14]. Cheong, B. C. (2024). Transparency and accountability in AI systems: safeguarding wellbeing in the age of algorithmic decision-making. *Frontiers in Human Dynamics*, 6.
[15]. Cheong, I., Caliskan, A., & Kohno, T. (2024). Safeguarding human values: rethinking US law for generative AI's societal impacts. *AI and Ethics*, 5(2), 1433–1459.
[16]. Eitel-Porter, R. (2020). Beyond the promise: implementing ethical AI. *AI and Ethics*, 1(1), 73–80.
[17]. Falco, G., Shneiderman, B., Badger, J., Carrier, R., Dahbura, A., Danks, D., Eling, M., Goodloe, A. E., Gupta, J., Hart, C., Jirotka, M., Johnson, H. H., LaPointe, C., Llorens, A. J., Mackworth, A. K., Maple, C., Pálsson, S. E., Pasquale, F., Winfield, A., & Yeong, Z. K. (2021). Governing AI safety through independent audits. *Nature Machine Intelligence*, 3(7), 566–571.
[18]. Griffin, T. A., Green, B. P., & Welie, J. V. M. (2023). The ethical agency of AI developers. *AI and Ethics*, 4(2), 179–188.
[19]. Hagendorff, T. (2020). The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds and Machines*, 30(1), 99–120.

ISSN :2584-2706

[20]. Häußermann, J. J., & Lütge, C. (2021). Community-in-the-loop: towards pluralistic value creation in AI, or—why AI needs business ethics. *AI and Ethics*, 2(2), 341–362.
[21]. Hickok, M. (2020). Lessons learned from AI ethics principles for future actions. *AI and Ethics*, 1(1), 41–47.
[22]. Ibitoye, A., Nkwo, M., & Orji, R. (2025). Rethinking responsible AI from ethical pillars to sociotechnical practice. *AI and Ethics*, 5(6), 6207–6223.
[23]. Irhas, I. (2026). Operationalizing AI Ethics: A Thematic Synthesis for Adaptive and Layered Governance. *MASALIQ*, 6(1), 411–428.
[24]. Jedličková, A. (2024). Ethical approaches in designing autonomous and intelligent systems: a comprehensive survey towards responsible development. *AI & Society*, 40(4), 2703–2716.
[25]. Krijger, J., Thuis, T., Ruiter, M. de, Ligthart, E., & Broekman, I. (2022). The AI ethics maturity model: a holistic approach to advancing ethical data science in organizations. *AI and Ethics*, 3(2), 355–367.
[26]. Lu, Q., Zhu, L., Xu, X., Whittle, J., Zowghi, D., & Jacquet, A. (2023). Responsible AI Pattern Catalogue: A Collection of Best Practices for AI Governance and Engineering. *ACM Computing Surveys*, 56(7), 1–35.
[27]. Minkkinen, M., Niukkanen, A., & Mäntymäki, M. (2022). What about investors? ESG analyses as tools for ethics-based AI auditing. *AI & Society*, 39(1), 329–343.
[28]. Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507.
[29]. Mökander, J. (2023). Auditing of AI: Legal, Ethical and Technical Approaches. *Digital Society*, 2(3).
[30]. Mökander, J., Axente, M., Casolari, F., & Floridi, L. (2021). Conformity Assessments and Post-market Monitoring: A Guide to the Role of Auditing in the Proposed European AI Regulation. *Minds and Machines*, 32(2), 241–268.
[31]. Mökander, J., & Floridi, L. (2022). Operationalising AI governance through ethics-based auditing: an industry case study. *AI and Ethics*, 3(2), 451–468.
[32]. Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2024). Auditing large language models: a three-layered approach. *AI and Ethics*, 4(4), 1085–1115.
[33]. Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2019). From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices. *Science and Engineering Ethics*, 26(4), 2141–2168.
[34]. Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2021). Operationalizing AI ethics: barriers, enablers and next steps. *AI & Society*, 38(1), 411–423.
[35]. Raji, I. D., Xu, P., Honigsberg, C., & Ho, D. E. (2022). Outsider Oversight: Designing a Third-Party Audit Ecosystem for AI Governance. *ACM AIES*, 557–571.