A Preliminary Investigation of a Methodology for Technically Informed Drafting and Revision of Technology Regulations

Ida Gioia Cristiano[1,†], Marco Angelini[1,*,†]

[1]Link University of Rome, Rome, Italy

Abstract

The rapid spread of technology, while beneficial, presents significant challenges, including the rise of disinformation and data security risks. Governmental responses via regulation are often criticized for lacking operational utility, misaligning with technical standards, and contributing to over-regulation. This paper proposes a methodological framework to support the revision and drafting of technology regulations, particularly for Artificial Intelligence (AI) and cybersecurity. The core method involves a structured conceptual analysis that maps a technology's technical latent space (operational requirements, risk conditions) to the corresponding regulatory latent space (e.g., EU AI Act, NIS2 Directive). Comparing these spaces allows for objective identification of misalignments: (1) regulatory gaps, (2) over-regulation, and (3) misaligned applicability thresholds. The framework integrates computational tools, such as Natural Language Processing, to provide transparent and reproducible analyses for regulators. Its goal is to improve the verifiability, coherence, and practical effectiveness of regulations by bridging the communication gap between legal and technical experts, ensuring legal obligations are grounded in technical reality.

Keywords

LaTeX class, paper template, paper formatting, CEUR-WS

1. Introduction

Over the last thirty years, the world has witnessed the rapid rise of technology at an unprecedented pace. Along with undeniable benefits, the spread of technology also came with potential drawbacks, such as the spread of disinformation, threats to human rights, and risks to personal data. International governments tried to cope with these problems by issuing regulations aimed at governing these technologies. Unfortunately, many of these regulations, primarily driven by jurisprudence, have been criticized as being ineffective, lacking specific operational guidelines, and misaligned with technical standards, thereby reinforcing the perception of over-regulation.

This paper advocates for the development of a methodological framework to support the revision and drafting of technology-related regulations (e.g., Artificial Intelligence, Cybersecurity). The proposed approach is based on a structured conceptual analysis that maps the **technical latent space** of a technology (operational principles, functional requirements, risk conditions) to the respective **regulatory latent space** defined in major European legislative instruments at the national and international levels, where applicable, like the EU AI Act and the NIS2 Directive.

The comparison between these two spaces enables the identification of three categories of normatively relevant misalignment: *(1) regulatory gaps*, where legal provisions fail to capture essential technical aspects making the regulation difficult to apply or subject to heterogeneity in interpretation; *(2) over-regulation*, where obligations are not grounded in realistic technical or operational conditions making the regulations difficult to digest and diluting its prescriptive nature toward its technical application; and *(3) misaligned applicability thresholds*, governing the previous two areas and allowing, also from technically

Joint National Conference on Cybersecurity (ITASEC SERICS 2026), February 09-13, 2026, Cagliari, IT

*Corresponding author.

†These authors contributed equally.

idagioia.cristiano@linkstudents.it (I. G. Cristiano); m.angelini@unilink.it (M. Angelini)

0000-0001-9051-6972 (M. Angelini)

perspective, to balance the alignment and avoid too broad or too narrow, producing uncertainty and implementation challenges.

Building on this conceptual model, the framework should integrate computational components, including Natural Language Processing techniques, semantic similarity models, and assisted regulatory-audit procedures. These tools should not replace human judgment, but provide transparent and reproducible analyses of the alignment between legal requirements and technical realities, highlighting consistencies, inconsistencies, and structural weaknesses within regulatory texts. Its target should be mainly regulators, allowing for a more objective identification of this alignment, supporting the identification of gaps during the process of creating and revising regulations, and serving as a means to enable regulators to communicate with technical experts.

Research opportunities arising from this model can be identified in the conceptual definition of the latent spaces, in terms of their entities, concepts, and relations that connect technical aspects to regulatory aspects in this specific direction. This includes the definition of decision-support platforms for policymakers and operators in AI and cybersecurity domains. The scientific contribution is twofold: it discusses a conceptual method for linking technical specifications and constraints to regulatory provisions, thereby enabling an operational foundation for quantifying the alignment, gaps, verifiability, coherence, and practical effectiveness of drafted and revised regulations in technological sectors. An example of its application is discussed in the context of the EU AI Act, and a subsequent discussion on research directions and challenges supports the advocated position.

2. Related Work

The rapid diffusion of Artificial Intelligence (AI) systems across social, economic and institutional contexts has generated a broad and articulated body of scientific literature addressing their technical foundations, governance implications, and regulatory challenges. AI is used in this section as an example of an emerging technical subject for which there is a rapidly increasing number of emergent regulations, well-representing the general idea of matching technology with regulations. Research examining the interaction between AI systems and legal norms has evolved along multiple, partially overlapping directions, reflecting the intrinsically socio-technical nature of these technologies. Within this landscape, existing contributions address the translation of legal principles into implementable technical requirements [1, 2], the development of methodologies for algorithmic auditing and system documentation to support accountability, the definition of technical standards for risk management and security [3], and the use of computational tools to assist legal and regulatory analysis [4]. More recently, the emergence of general-purpose AI systems has fostered increasing attention to adaptive governance models and human oversight mechanisms.

2.1. Legal Principles, Technical Requirements, and Computational Methods

A first research strand investigates the conceptual distance between abstract legal principles and their practical implementation in AI systems. Selbst et al. show that translating notions such as fairness or accountability into measurable technical metrics requires careful management of abstraction levels: overly general principles cannot be directly implemented, while excessively specific operationalizations fail to capture socio-technical complexity [5]. Similarly, Wachter, Mittelstadt, and Floridi demonstrate that even seemingly clear legal expectations, such as the so-called “right to explanation” cannot be meaningfully operationalized when existing machine-learning architectures do not support adequate forms of transparency [6].

Further contributions emphasize that algorithmic transparency and accountability cannot be achieved solely through normative prescriptions but require governance frameworks capable of linking system design, documentation, and evaluation phases. In this perspective, traceability-oriented approaches emphasize the need to integrate legal requirements into technical and organizational processes throughout the AI lifecycle [7]. Taken together, these studies share an analytical orientation that moves from legal concepts toward their translation into technical constraints.

Closely related to this line of work, a growing body of research explores the use of computational methods, particularly Natural Language Processing (NLP) techniques, to support legal and regulatory analysis. Approaches focused on data documentation highlight the necessity of formalizing dataset quality and representativeness [8]. Large-scale surveys on deepfake manipulation and detection further illustrate the difficulty of evaluating biometric systems and the fragility of current detection mechanisms [9]. More broadly, empirical studies demonstrate that NLP techniques and large language models can facilitate the classification of regulatory texts, the extraction of legal obligations, and the identification of inconsistencies, suggesting their potential as assistive tools for regulators and auditors.

2.2. Algorithmic Auditing, Documentation, and Risk Management

A second research strand focuses on tools and methodologies for validating AI systems, documenting their behavior, and enabling accountability across the system lifecycle. Mitchell et al. introduce *Model Cards*, proposing standardized documentation formats for reporting model performance, limitations, and intended use cases [10]. Complementarily, Gebru et al. propose *Datasheets for Datasets*, emphasizing data provenance, quality assessment, and context-dependent representativeness [8]. Raji et al. identify an *accountability gap* arising from the absence of structured documentation practices throughout the AI lifecycle [11].

Recent work extends these auditing paradigms to large language models (LLMs). Mökander et al. propose a three-layered audit framework that considers model behavior, deployment context, and governance structures [12], while Amirizaniani et al. introduce *LLMAuditor*, a human-in-the-loop approach to auditing LLMs [13]. Across this strand, the literature consistently argues that auditing practices cannot be effective in the absence of structured documentation standards, while regulatory texts rarely specify which standards should be adopted.

Closely connected to auditing practices is a substantial body of work on AI risk management. The *NIST AI Risk Management Framework* provides operational tools for identifying, assessing, and mitigating risks, including robustness measures, drift detection, uncertainty quantification, and attack-surface analysis [14]. Technical research on adversarial robustness and model uncertainty further highlights the fragility of machine-learning systems in real-world settings, showing how models fail under adversarial perturbations [15], evade detection mechanisms [16], or produce unreliable outputs due to calibration errors [17]. Collectively, these contributions demonstrate that technical standards for safety and robustness exist, yet they are not systematically embedded within legal and regulatory obligations. It is still not investigated how to consider them during the drafting of new regulations in an objective and quantifiable manner.

2.3. Adaptive Governance and Human Oversight

The increasing deployment of general-purpose AI systems and LLMs has intensified interest in adaptive governance models and human oversight mechanisms. Human-in-the-loop auditing strategies show that meaningful oversight requires well-defined intervention criteria, escalation procedures, and adequate operator competence [12]. This literature consistently reveals a gap between high-level regulatory expectations, such as “effective human oversight”, and the technical realities of complex and dynamic AI pipelines.

Overall, the research strands reviewed here primarily address the relationship between technology and regulation by starting from legal texts and seeking to translate them into technical requirements or compliance obligations. Less attention has been devoted to perspectives that use technical governance artifacts and operational practices as analytical lenses to evaluate the coherence, applicability, and effectiveness of regulatory provisions. It is within this methodological space that the present work is positioned.

3. A conceptual space to match technological and regulatory concepts

3.1. Method Description

The hypothesized conceptual model is conceived as a structured analytical framework for examining technology-oriented regulation, explicitly assuming technical aspects as the starting point of normative analysis. In technology-intensive domains, legal prescriptions cannot be adequately understood or evaluated in isolation from their broader context. Rather, they must be interpreted in light of the technical, operational, and organizational instruments that enable their concrete implementation and simultaneously delineate their limits of applicability. In this perspective, the model should adopt an explicitly technology-informed orientation, whereby regulatory provisions are not treated as primary inputs but as normative outcomes that incorporate, either explicitly or implicitly, assumptions concerning technical feasibility, measurability of requirements, verifiability of obligations, risk manageability, and the possibility of effective human control and oversight.

The model, therefore, rejects a purely abstract or doctrinal reading of legal texts and instead grounds regulatory analysis in the operational reality of technological governance. Technical instruments are not conceived as merely executory tools subordinate to legal norms, but as epistemic and organizational artifacts that actively shape regulatory design, interpretation, and enforcement. By systematically examining these instruments, the framework aims to make explicit which technical capabilities are realistically available to support regulatory objectives and to what extent such capabilities are accurately reflected in legal provisions. In this sense, the methodology is intended as a support tool for regulatory design and revision, enabling legislators and regulators to align normative ambitions with technical realities.

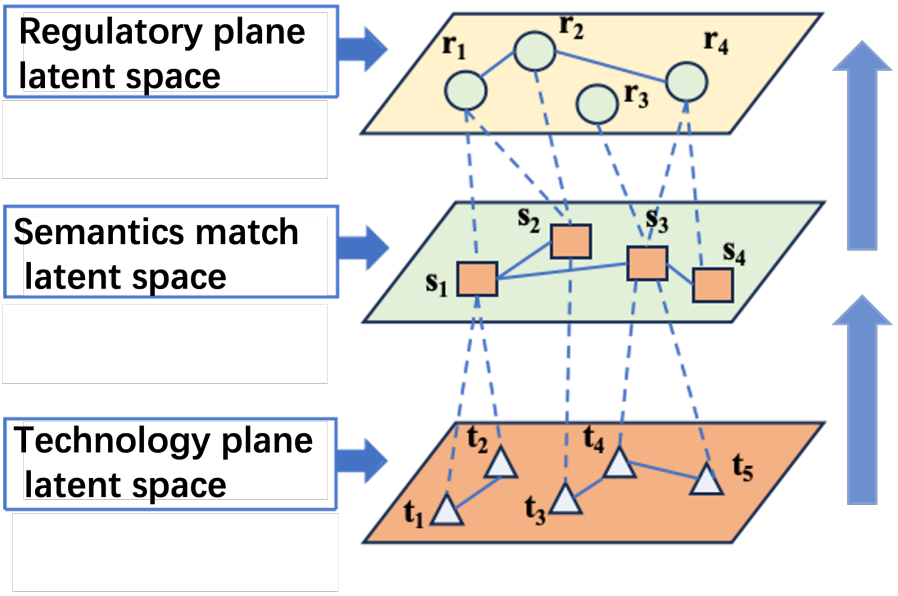

Figure 1: The conceptual space matching and informing regulatory concepts by technology specifications

Figure 1 provides an overview of the proposed model. In this conceptual view, taxonomies of general concepts related to technology and its regulations (for a generic technology) are derived from the literature and form cardinal points of two latent spaces. These two planes (technological and regulatory, where the first must inform the second) are not directly connected, but mediated by a third plane, with the goal of matching the semantics of technology constraints and specification to the mandates and prescriptions of a regulation, informing them. This match may have different meanings depending on the kind of relation established for it (e.g., a technical prescription informing an article, providing an exception or corner case to consider, or influencing the description of specific technical aspects among other cases).

3.2. Technology Plane

Within the model, technical instruments constitute the primary analytical lens through which regulatory provisions are interpreted and assessed. These instruments correspond to the set of practices, procedures, and artifacts effectively employed in the governance of technological systems. Policies provide an overarching governance framework by defining objectives, accountability principles, and acceptable levels of risk. As emphasized in the literature on ethical and responsible AI governance, policies play a central role in translating abstract values into operational priorities and risk-based decision-making structures [18, 19]. They inform regulatory interpretation by clarifying how high-level principles such as transparency or responsibility should be operationalized in concrete contexts.

Procedures operationalize these governance objectives by translating policy-level commitments into repeatable practices applied across the system lifecycle, including design, development, deployment, and monitoring. From a regulatory perspective, procedures reveal whether legal obligations can be realistically implemented in a consistent and auditable manner, thus exposing potential gaps between normative intent and operational feasibility [20]. Controls, both technical and organizational, enable ongoing supervision and risk mitigation by constraining system behavior and detecting deviations from predefined objectives. As highlighted in studies on audit and compliance, controls function as the primary interface between regulatory expectations and operational enforcement [21, 22].

Documentation structures information concerning systems, models, and datasets, recording decisions, assumptions, and modifications over time. In regulatory settings, documentation is widely recognized as a foundational instrument for accountability, auditability, and ex post scrutiny, particularly in complex socio-technical systems [23, 24]. Verification steps further operationalize governance by transforming abstract regulatory concepts, such as robustness or accountability, into assessable evidence that can support conformity assessments and regulatory audits [7]. Applicability considerations define the conditions under which specific technical measures are activated, taking into account contextual factors such as system function, deployment scale, and risk profile. These considerations are crucial for avoiding overgeneralization and ensuring proportionality in regulatory application.

3.3. Regulatory Plane

The regulatory plane of the model encompasses the heterogeneous set of normative elements through which legal systems articulate expectations for technological systems. Regulatory instruments range from general principles to operative provisions structured into articles, paragraphs, and clauses, as well as exceptions, technical annexes, and references to external standards or normative corpora. Legal theory distinguishes principles as norms that express optimization requirements and guide interpretation, while rules and articles translate these principles into concrete obligations, prohibitions, or permissions [25, 26]. Paragraphs and clauses further specify the scope and conditions of application, whereas exceptions modulate normative reach to address contextual variability.

Technical annexes and references to standards are particularly relevant in technology-oriented regulation, as they seek to operationalize legal requirements by incorporating technical specifications or recognized best practices. However, as scholarship on principles-based and risk-based regulation has shown, the effectiveness of such instruments depends critically on their alignment with existing technical practices and verification capabilities [22, 27]. Regulatory texts, therefore, implicitly rely on assumptions about the availability and maturity of technical instruments that can support compliance and enforcement.

3.4. Insights and Analysis Results

The quantitative modeling of the interplay between technical plane concepts and regulatory plane concepts defines what can be conceptualized as a shared analytical space, in which alignment and misalignment patterns emerge in a quantifiable manner, making it useful for the activity of drafting and revising regulations. In particular, this method allows for the identification and quantification of four main areas, summarized in Figure 2.

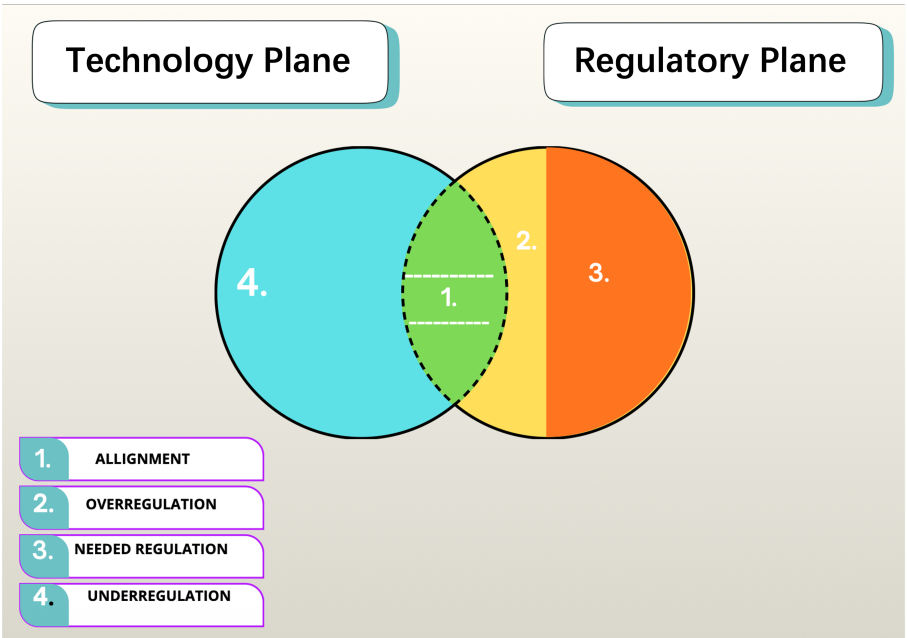

Figure 2: Conceptual diagram illustrating the relation between the Technology plane latent Space and the Regulatory plane latent Space.

By mapping technical instruments to regulatory elements, the methodology enables the identification of situations in which regulatory provisions are fully supported by existing technical practices, resulting in substantive alignment between normative objectives and operational capabilities (*alignment area*, see Figure 2 area 1 in green). In this scenario, the regulations capture important technical aspects of the technology. In other cases, the analysis reveals forms of disalignment between the two planes: the area where regulation provisions lack sufficient operational specification to be meaningfully implemented or verified through available technical instruments can be identified as *underregulation* (see Figure 2 area 4 in azure). Conversely, instances of *overregulation* emerge when legal requirements presuppose levels of technical control, certainty, or measurability that exceed current technological capabilities, thereby generating feasibility and verification gaps (see Figure 2 area 2 in yellow). Both these areas should be minimized and at least monitored to increase the efficacy and effectiveness of the regulations, and our proposal enables their identification.

Finally, the proposed framework identifies another important area where alignment is not present by default: the area of *needed regulation* (see Figure 2 area 3 in orange). It represents normative elements that must necessarily be present in a regulation due to its harmonization with foundational legal principles and existing linked regulations. Such elements may include, for example, accountability structures and minimum documentation and oversight requirements, which are consistently recognized in the literature on algorithmic and technological governance [28, 29].

Through this structured analysis, the proposed methodology does not aim to automate compliance assessment or replace legal judgment. Rather, it provides a transparent and reproducible analytical basis for evaluating how technical realities inform, constrain, or challenge regulatory design, thereby supporting legislators and regulators in refining technology-oriented legal frameworks providing a support for them.

As depicted in Figure 2, the Technology Plane is designed to inform the Regulatory Plane, rather than being subordinate to it. This directional relation is fundamental to the proposed methodology, as the model explicitly assumes that technical feasibility and measurability should inform normative design.

4. Case Study: Application to the EU AI Act

This section presents a demonstration case study on the EU AI Act, conducted using a semi-assisted application of the proposed method. The method leverages BERTopic and a local version of GPT-OSS for representing the Technical and Regulatory latent space, and involves two human experts for reconciling the analysis and identifying gaps. This approach enables the analysis of how policies, operational procedures, controls, documentation, verification steps, applicability considerations, and other relevant technical instruments inform principles, articles, clauses, exceptions, and technical annexes of the EU AI Act. Results are reported in Table 1.

Technical instruments are conceived as structured socio-technical artifacts encompassing operational assumptions, applicability constraints, epistemic limits, and auditing and verification tools. They provide operational content to regulatory prescriptions, translating abstract concepts such as transparency, accountability, or robustness into verifiable requirements and standard procedures [30, 7, 10, 8]. Similarly, regulatory instruments include principles, articles, clauses, exceptions, technical annexes, and references to standards and guidelines; they define the regulatory framework within which technical instruments must operate.

Within the **Technical Plane**, the term "policy" refers to formal governance documents defining objectives, accountability structures, and acceptable risk levels. "Procedures" indicate repeatable operational workflows applied throughout the system lifecycle. "Controls" encompass both technical and organizational mechanisms used to constrain system behavior and mitigate risks. "Documentation" includes structured artifacts such as model cards, audit logs, and risk registers. "Verification steps" correspond to measurable validation mechanisms, such as audits, testing procedures, or conformity assessments.

Within the **Regulatory Plane**, "principles" represent high-level normative objectives, while "articles" and "clauses" translate these objectives into concrete obligations. "Exceptions" limit or contextualize normative scope. "Technical annexes" and "references to standards" incorporate external technical specifications into the legal framework.

The application of the proposed methodological model enables a systematic comparison between the **technical latent space** and the **regulatory latent space**, highlighting recurring patterns of misalignment between available technical instruments and regulatory prescriptions.

To clarify how the proposed methodology operates, Article 23 is analyzed through a structured triple:

1. **Legal Aspect** Article 23 requires essential and important entities to submit an early warning within 24 hours, followed by a detailed notification within 72 hours, and a final report. The legal prescription assumes timely detection and classification of significant incidents.
2. **Technical Aspect** From a cybersecurity perspective, incident detection depends on the presence of Security Operations Centers (SOC), SIEM systems, detection engineering capabilities, and predefined incident classification procedures. Empirical studies show that Mean Time To Detect (MTTD) often exceeds 24 hours, especially in organizations with limited monitoring maturity.
3. **Relational Pattern** The comparison between the legal and technical aspects reveals a feasibility gap. The regulation presupposes detection capabilities that are not uniformly available across entities. The obligation is normatively clear but technically dependent on maturity levels not explicitly addressed in the Directive.

This triple-structure exemplifies how the methodology operationalizes the comparison between the two latent spaces.

4.1. Results

The analysis of grouped articles shows how specific classes of technical instruments support particular categories of regulatory requirements. For instance, documentation and transparency obligations (Articles 13, 69–71) are underpinned by model cards and datasheets, enabling verification of traceability, model characteristics, and dataset limitations [10, 8]. In these cases, **format gaps** and **procedural**

Table 1
Mapping between Technical and Regulatory Planes

Technical Plane	Regulatory Plane	Mapping	Emerging Gap	Articles
Policy	General principles and articles	Policies define governance objectives and risk levels, clarifying how to interpret transparency and accountability principles	Alignment	3, 10–13
Documentation	Articles and technical annexes	Technical documentation (model cards, datasheets) provide verifiable evidence of compliance with transparency requirements	Aligned	10–13
Procedures	Operational clauses and articles	Procedures translate principles into operational practices, indicating how to implement security and robustness requirements	Aligned	25–28
Controls	Technical annexes	Technical and organizational controls enable verification of compliance with operational requirements specified in annexes	Aligned	29–31
Verification steps	Articles, clauses	Verification steps transform abstract requirements into concrete evidence, allowing validation of compliance	Aligned	32–35
Procedures	Clauses and articles	Procedures show that some regulatory provisions lack sufficient operational guidance, indicating gaps	Underregulation	1–2, 4, 6–7, 14
Controls	Articles	Controls show that monitoring requirements are not fully supported by technical instruments	Underregulation	36–40
Procedures and controls	Articles and clauses	Procedures and controls highlight that regulatory requirements demand technical capabilities not yet consolidated	Overregulation	5, 8, 15, 52
Policy	General principles	Policies highlight fundamental elements that must be present regardless of context, such as non-derogable principles	Needed Regulation	21–24
Documentation and controls	Annexes and normative references	Technical documentation guides the interpretation of annexes and references, ensuring minimum necessary compliance	Needed Regulation	Tech. annexes and references

underspecification emerge because the regulation requires documentation without specifying standard formats or explicit verification protocols.

Provisions requiring robustness, accuracy, resilience, and security (Articles 9, 15) are informed by technical risk assessment tools, robustness metrics, vulnerability analyses, and monitoring frameworks such as NIST AI RMF 1.0 and ISO/IEC 23894:2023 [14, 15, 16, 17]. The analysis reveals **feasibility gaps** and **technical underspecification**, highlighting how regulatory requirements impose security conditions that cannot always be verified using current methodologies. For provisions addressing emerging risks and prohibited practices (Articles 5, 52), the literature on deepfakes, behavioral manipulation, and LLM auditing identifies **measurement gaps** and **applicability issues**, as the regulation presupposes detection and control capabilities that are not yet technically consolidated.

Operational procedures and controls outlined in Articles 6–7, 53–55, and 69–71 are mapped to auditing, human oversight, and regulatory sandbox instruments. The analysis highlights **protocol gaps**, where the regulation defines objectives or institutional mechanisms without providing operational technical guidelines, leaving wide discretion in implementation and compliance evaluation.

The proposed methodological approach enables systematic mapping of each article to its corresponding set of technical instruments, identifying key gaps: **over-regulation**, when regulatory prescriptions require unrealistic capabilities; **technical underspecification**, when concrete metrics or procedures are missing; **normative vagueness**, when abstract concepts lack technical correspondence; and **applicability and verification gaps**, when practical implementation is uncertain or depends on external conditions. These recurring patterns emerge consistently across the analyzed groupings, allowing an integrated and systematic reading of technical-regulatory misalignments.

This analysis directly supports legislators and regulatory stakeholders by providing concrete evidence of where regulatory prescriptions can be strengthened through standardization, procedural specification, or additional technical support. Consequently, the methodology facilitates the design of more robust governance instruments, harmonizing regulatory needs and technical constraints, and proposes operational guidelines for audits, regulatory sandboxes, and technical compliance.

5. Discussion

The analysis conducted in this work illustrates how the current regulatory approach to Artificial Intelligence in the EU faces structural challenges when confronted with the technical properties of modern AI systems. While the AI Act represents one of the most advanced attempts to regulate complex, adaptive, and partially opaque technologies, the case study reveals several forms of misalignment between legal requirements and technical feasibility.

A first key finding concerns the persistent gap between high-level normative principles and the operational metrics needed to verify compliance. Requirements such as transparency, robustness, fairness, and effective human oversight are conceptually well-defined in the regulation, but lack corresponding technical indicators or standardized procedures. As a consequence, compliance may risk becoming formalistic, based on documentation rather than on measurable technical performance.

Second, the analysis highlights that many regulatory obligations presuppose technical conditions that are not yet achievable for state-of-the-art AI systems. Examples include complete traceability of training data, ongoing explainability of complex models, and reliable detection of manipulative behaviors or deepfakes. These obligations reflect desirable forms of governance, but require the development of new methodologies, standards, and evaluation frameworks before they can be consistently implemented.

Third, the dual-space approach used in this paper demonstrates the importance of systematically mapping regulatory prescriptions to specific technical capabilities and constraints. This type of structured comparison enables policymakers and regulators to identify where obligations are realistic, where additional technical tools are necessary, and where regulatory expectations need to be refined to prevent over-regulation or inconsistent enforcement across Member States. It also supports the design of regulatory sandboxes and testing environments that incorporate meaningful technical evaluations rather than purely administrative checks.

Moreover, the discussion highlights the relevance of integrating computational tools, such as NLP-based analysis, semantic similarity models, and automated auditing pipelines, into the regulatory process. These tools cannot replace human judgment, but they can provide systematic and replicable insights into regulatory coherence, coverage, and feasibility. Their adoption can also enhance transparency by making explicit the assumptions and boundaries of technical assessments.

Ultimately, the findings suggest that effective AI governance necessitates ongoing dialogue between technical and legal communities, as well as the progressive development of harmonized standards for documentation, testing, and risk assessment. As AI systems evolve rapidly, regulatory frameworks must incorporate mechanisms for iterative refinement and technical updating to remain aligned with emerging capabilities and risks. The methodological model proposed in this work contributes to this goal by offering a structured foundation for future regulatory innovation.

5.0.1. Research Challenges and Implications

One of the main research challenges emerging from this investigation concerns the operationalization and measurement of alignment between the technical plane and the regulatory plane latent spaces. From a conceptual perspective, measuring alignment requires the construction of semantic metrics that can compare the operational definitions embedded in policies, procedures, and technical documentation with the normative content expressed in principles, articles, and clauses. This entails the development of models capable of capturing not only terminological correspondence but also functional coherence between regulatory objectives and technical modes of implementation, moving beyond a purely formal notion of compliance and requiring data from human interpretation and past experiences in this exercise.

A further challenge relates to the semantic overlap and differences between heterogeneous languages. Technical jargon, oriented toward processes, metrics, and operational instruments, and legal jargon, grounded in principles, obligations, and abstract categories, follow different expressive forms and purposes. The proposed model provides a foundation that needs to be better tailored to these more subtle differences and capable of providing a conceptual structure that can capture not only overlaps but also implications. In this perspective, desired characteristics can be identified and made explicit by highlighting the technical assumptions embedded in regulatory choices, as well as the legal constraints that shape technical design and technology governance practices.

Considering policymakers, a research challenge is to provide support during the drafting and revision of technology-related regulations. By analyzing the interaction between technical instruments and regulatory components, the approach enables a more precise identification of the level of specification required in articles, clauses, and technical annexes, as well as of the regulatory terminology that can be effectively grounded in existing technical practices. Looking forward, research efforts should focus on decision-support or automated recommendation systems specifically designed for policymakers, capable of deriving and suggesting regulatory formulations that are consistent with available technical capabilities and providing established verification mechanisms. While some efforts are beginning to emerge in this direction [31], they are still focused on the pure legal aspects, with more attention devoted to the interplay with technological aspects.

Finally, this work opens several methodological and computational research directions. The integration of Natural Language Processing techniques and semantic analysis could enable the automation of identifying patterns of alignment, underregulation, and overregulation, facilitating the large-scale analysis of complex regulatory texts. From this standpoint, the proposed model should be understood not as a final solution, but as a conceptual foundation for the development of more adaptive governance tools, capable of evolving alongside the technologies they are intended to regulate.

6. Conclusion

This paper has examined the relationship between the technical specifications and constraints of a Technology and the related regulatory requirements through a conceptual dual-space framework, with the goal of providing quantitative support for the degree of misalignment to inform the drafting and

revisions of the regulation. By modeling the regulatory plane and the technical plane latent spaces, the proposed model provides a foundation for identifying regulatory gaps, over-regulation, and ambiguous applicability thresholds within a technology-related regulation. The preliminary mapping of these gaps for the EU AI Act performed in the case study highlights where legal provisions are aligned with current technological capabilities and where they rely on assumptions that are not yet technically feasible, misaligned, and under- or over-regulated.

The identification of computational tools, such as NLP-based analysis and semantic similarity models, may further strengthen regulatory assessment processes during the drafting and revision stages. These tools enable transparent, repeatable, and scalable evaluations of legal texts, allowing for the detection of inconsistencies, redundancies, and areas requiring clarification. They are not intended to replace the interpretive role of policymakers, but rather support it by providing structured evidence about the technical feasibility of regulatory obligations.

The initial findings of this work, which focused on generative AI, suggest that effective technology regulations necessitate ongoing collaboration between technical expertise and legal reasoning, as well as the progressive development of harmonized standards for documentation, testing, and risk management. As AI systems continue to evolve in complexity, regulatory frameworks must incorporate mechanisms for iterative updating and evidence-based refinement to remain coherent, applicable, and enforceable.

The conceptual model advocated in this paper contributed to this objective by offering an operational approach to linking legal provisions with technical realities, even in its initial form. Its full implementation may provide a foundation for future research on assisted regulatory drafting, analysis, and revisions, support to policymakers in designing more technically grounded regulations, and allow for the correct quantification of under- and over-regulation, allowing for a more objective specification of these elements and improving the applicability and public perception of these regulations, providing quantifiable evidence. These directions are identified as future research efforts.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] M. Angelini, C. Ciccotelli, L. Franchina, A. Marchetti-Spaccamela, L. Querzoni, Italian national framework for cybersecurity and data protection, in: L. Antunes, M. Naldi, G. F. Italiano, K. Rannenberg, P. Drogkaris (Eds.), Privacy Technologies and Policy, Springer International Publishing, Cham, 2020, pp. 127–142.

[2] T. Sammarone, M. Angelini, A human-ai assisted methodology for the generation of a cybersecurity controls library, in: 2025 IEEE International Conference on Cyber Humanities (IEEE-CH), 2025, pp. 1–7. doi:10.1109/IEEE-CH65308.2025.11279411.

[3] M. P. Carello, A. Marchetti-Spaccamela, L. Querzoni, M. Angelini, Sok: Cybersecurity regulations, standards and guidelines for the healthcare sector *, in: 2023 IEEE International Conference on Intelligence and Security Informatics (ISI), 2023, pp. 1–6. doi:10.1109/ISI58743.2023.10297246.

[4] M. Angelini, S. Lenti, G. Santucci, Crumbs: A cyber security framework browser, in: 2017 IEEE Symposium on Visualization for Cyber Security (VizSec), 2017, pp. 1–8. doi:10.1109/VIZSEC.2017.8062194.

[5] A. D. Selbst, D. Boyd, S. A. Friedler, S. Venkatasubramanian, J. Vertesi, Fairness and abstraction in sociotechnical systems, in: Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*), 2019, pp. 59–68. doi:10.1145/3287560.3287598.

[6] S. Wachter, B. Mittelstadt, L. Floridi, Why a right to explanation does not exist under the gdpr, International Data Privacy Law 7 (2017) 76–99. doi:10.1093/idpl/ipx005.

[7] J. A. Kroll, Outlining traceability: A principles-based approach to algorithmic accountability, Minds and Machines (2021). doi:10.1007/s11023-020-09542-4.

[8] T. Gebru, et al., Datasheets for datasets, 2018. arXiv:1803.09010.
[9] Y. Mirsky, W. Lee, The creation and detection of deepfakes: A survey, ACM Computing Surveys 54 (2021) 1–41. doi:10.1145/3425780.
[10] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, T. Gebru, Model cards for model reporting, in: Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*), 2019, pp. 220–229. doi:10.1145/3287560.3287596.
[11] I. D. Raji, et al., Closing the ai accountability gap, Proceedings of the ACM Conference on Fairness, Accountability, and Transparency (FAccT) (2020). doi:10.1145/3351095.3372873.
[12] J. Mökander, et al., Auditing large language models: A three-layered approach, 2023. arXiv:2303.04430.
[13] P. Amirizaniani, et al., Llmauditor: A human-in-the-loop framework for auditing llms, 2024. arXiv:2402.05099.
[14] National Institute of Standards and Technology (NIST), Artificial intelligence risk management framework (ai rmf 1.0), 2023.
[15] A. Madry, et al., Towards deep learning models resistant to adversarial attacks, in: International Conference on Learning Representations (ICLR), 2018.
[16] N. Carlini, D. Wagner, Adversarial examples are not easily detected, arXiv preprint arXiv:1705.07263 (2017).
[17] D. Hendrycks, et al., Uncertainty and robustness in deep learning, arXiv preprint arXiv:2102.12192 (2021).
[18] L. Floridi, J. Cowls, M. Beltrametti, et al., Ai4people—an ethical framework for a good ai society, Minds and Machines 28 (2018) 689–707. doi:10.1007/s11023-018-9482-5.
[19] D. Kaye, Governing artificial intelligence through human rights, European Journal of International Law 31 (2020) 649–676. doi:10.1093/ejil/chaa038.
[20] R. Baldwin, J. Black, Driving priorities in risk-based regulation: What's the problem?, Journal of Law and Society 39 (2012) 542–575. doi:10.1111/j.1467-6478.2012.00598.x.
[21] M. Power, The Audit Society: Rituals of Verification, Oxford University Press, 1997.
[22] J. Black, Forms and paradoxes of principles-based regulation, Capital Markets Law Journal 3 (2008) 425–457. doi:10.1093/cmlj/kmn017.
[23] K. A. Bamberger, D. K. Mulligan, Privacy on the books and on the ground, Stanford Law Review 63 (2010) 247–315.
[24] A. A. Winecoff, M. Bogen, Improving governance outcomes through ai documentation: Bridging theory and practice, arXiv preprint arXiv:2409.08960 (2024). URL: https://arxiv.org/abs/2409.08960.
[25] R. Dworkin, Taking Rights Seriously, Harvard University Press, 1977.
[26] R. Alexy, A Theory of Constitutional Rights, Oxford University Press, 2002.
[27] C. R. Sunstein, Problems with rules, California Law Review 83 (1995) 953–1026.
[28] K. Yeung, Algorithmic regulation: A critical interrogation, Regulation & Governance 12 (2018) 505–523. doi:10.1111/rego.12158.
[29] R. Brownsword, Law, technology and society: Re-imagining the regulatory environment, Law, Innovation and Technology 11 (2019) 1–32. doi:10.1080/17579961.2019.1574407.
[30] B. D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, L. Floridi, The ethics of algorithms: Mapping the debate, Big Data & Society 3 (2016) 2053951716679679. doi:10.1177/2053951716679679.
[31] D. Fürst, M. El-Assady, D. A. Keim, M. T. Fischer, Challenges and opportunities for visual analytics in jurisprudence, Artificial Intelligence and Law (2025) 1–32.