Stanford – Vienna
Transatlantic Technology Law Forum

A joint initiative of
Stanford Law School and the University of Vienna School of Law

TTLF Working Papers

No. 138

Unleashing Innovation, Leashing Responsibility: Risk, Responsibility, and Regulatory Form in Transatlantic AI Law

Umberto Nizza

2025

TTLF Working Papers

Editors: Siegfried Fina, Mark Lemley, and Roland Vogl

About the TTLF Working Papers

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at http://ttlf.stanford.edu. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum
http://ttlf.stanford.edu

Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

University of Vienna School of Law
Department of Business Law
Schottenbastei 10-16
1010 Vienna, Austria

About the Author

Umberto Nizza is an Adjunct Professor at the University of Bologna, Italy. He holds a Ph.D. in Law, Institutions, and Economics from Collegio Carlo Alberto, and has a multidisciplinary background in both law and economics. Before entering academia, he worked as an attorney and as an international consultant. His research focuses on artificial intelligence, judicial efficiency, and economic innovation. Nizza's work has appeared in leading legal and economic journals, and he has contributed to several EU-funded projects, including Horizon 2020, COST Action, and Jean Monnet Modules. He also serves as editor of The Innovation Letter within the Journal of Law, Market & Innovation. His current research explores how AI regulation shapes technological innovation, how emerging technologies transform the legal profession, and how regulatory frameworks can adapt to data-driven and rapidly evolving environments. Another strand of his work examines how cultural and institutional factors influence the economic behavior of judicial actors, employing advanced econometric methods. In addition, he studies political productivity, absenteeism, and rent-seeking behavior. Nizza welcomes opportunities for collaboration on future projects, contributing expertise at the intersection of law, economics, and technology, with the aim of informing evidence-based policy and governance.

General Note about the Content

The opinions expressed in this paper are those of the author and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

Suggested Citation

This TTLF Working Paper should be cited as:
Umberto Nizza, Unleashing Innovation, Leashing Responsibility: Risk, Responsibility, and Regulatory Form in Transatlantic AI Law, Stanford-Vienna TTLF Working Paper No. 138, http://ttlf.stanford.edu.

Copyright

Abstract

This paper conducts a comparative legal analysis of the European Union's and the United States' regulatory architectures for artificial intelligence (AI) through a computational methodology grounded in natural language processing (NLP). By applying techniques such as TF-IDF vectorization, word embedding, principal component analysis, and sentiment analysis, the study examines the semantic structure, lexical density, and normative rationalities embedded in the EU Artificial Intelligence Act (Regulation (EU) 2024/1689) and the recent U.S. Executive Orders on AI. The empirical findings reveal a profound structural and conceptual divergence between the two regulatory systems. The European corpus exhibits a high concentration of juridical and procedural terminology—high-risk, compliance, provider, surveillance, conformity—reflecting an ex ante, risk-preventive, and rights-protective legal rationality. Conversely, the U.S. corpus displays a more heterogeneous lexical distribution dominated by terms such as policy, national, promote, leadership, and agency, indicative of an executive-driven model of governance grounded in discretion, inter-agency coordination, and innovation enablement. Quantitative and semantic analyses confirm the existence of two distinct regulatory epistemologies: a European model characterized by legal formalism, normative determinacy, and precautionary proportionality; and a U.S. model distinguished by procedural agility, policy pragmatism, and strategic competitiveness. An additional interpretive layer employing large language models (ChatGPT, Copilot, Gemini, and Mistral) substantiates these results, demonstrating that LLMs are capable—under controlled conditions—of identifying latent regulatory tensions, coherence asymmetries, and potential avenues for transatlantic harmonization. At a theoretical level, the paper deconstructs what it terms the AI regulatory paradox: both jurisdictions pursue the same goals of safety, trust, and accountability, yet they do so through structurally incommensurable legal grammars. Building on this finding, the study advances a proposal for an adaptive, risk-based governance paradigm that integrates the European model's normative density and rights-based safeguards with the American system's procedural flexibility and innovation-oriented design. Beyond providing empirical evidence of transatlantic divergence, the paper demonstrates how NLP-based computational tools can serve as meta-regulatory instruments for the design, calibration, and critical assessment of future legal architectures in the age of algorithmic governance.

Table of content

1. Introduction

The unprecedented acceleration in the development and deployment of artificial intelligence (AI) systems has prompted regulators across the globe to deal with fundamental questions of governance, responsibility, and risk allocation[1]. Among the most prominent and ideologically contrasting approaches to AI regulation are those adopted by the European Union (EU) and the United States (US). The EU, driven by a tradition of normative constitutionalism and precautionary risk management[2], has proposed a comprehensive, binding, and risk-tiered legal framework: The Artificial Intelligence Act (AI Act)[3]. By contrast, the US has predominantly relied on executive orders[4], voluntary guidelines, and sector-specific norms, resulting in a fragmented and more complex regulatory landscape[5]. This transatlantic divergence is neither accidental nor purely structural; rather, it is emblematic of two fundamentally distinct legal paradigms. The EU regulatory approach rests on *ex ante* legal harmonization, operationalized through uniform, binding standards designed to address and mitigate identified categories of AI risk, while the US approach favors *ex post* administrative

[1] Cf., *inter alia*, Teemu Birkstedt, Matti Minkkinen, Anushree Tandon, and Matti Mäntymäki. 'AI governance: themes, knowledge gaps and future agendas' [2023] Internet Research 133; Virginia Dignum, 'Responsibility and artificial intelligence' in Markus Dirk Dubber, Frank Pasquale, and Sunit Das (eds), *The oxford handbook of ethics of AI* (Oxford University Press 2020); Tamás Fézer, 'Upside Down: Liability, Risk Allocation and Artificial Intelligence' [2024] Pro Publico Bono–Public Administration 85.

[2] See, in this regard, among others, Fisher, E. (2002). Precaution, precaution everywhere: developing a 'common understanding'of the precautionary principle in the European Community. *Maastricht Journal of European and Comparative Law*, 9(1), 7-28.

[3] For a comprehensive commentary on the Regulation (EU) 2024/1689, see Theodoros Karathanasis, *The EU Artificial Intelligence (AI) Act: A Commentary* (Wolters Kluwer 2024).

[4] For a comprehensive analysis of the *Executive order on the safe, secure, and trustworthy development and use of artificial intelligence*, issued during Biden's presidency – which is beyond the interest and scope of the present study – see*, ex multis*, Wörsdörfer, M. (2025). Biden's executive order on AI: Strengths, weaknesses, and possible reform steps. *AI and Ethics*, 5(2), 1669-1683.

[5] Evidence to this effect is provided by David Blumenthal and Amulya Marellapudi, 'More Fragmented, More Complex: State Regulation of AI in Health Care' [2025] NEJM AI 1.

coordination, wherein oversight and remediation are implemented reactively through agency discretion after issues have materialized[6].

While existing comparative scholarship has predominantly emphasized institutional divergences in legislation between the two sides of the Atlantic[7], this paper aims to go beyond a merely institutional comparison by delving into the underlying decision-making rationale and governance philosophy surrounding AI regulation. Employing natural language processing (NLP) techniques to systematically examine the textual architecture and semantic patterns of regulatory instruments we compare – applying, among others, TF-IDF vectorization, word frequency analysis, linguistic lemmatization, and large language models – the lexical priorities embedded in the EU AI Act and relevant US Executive Orders on AI. This form of computational textual analysis enables us to move beyond mere doctrinal interpretation and to uncover implicit regulatory rationalities encoded in language choices, discursive framing, and terminological emphasis. The results expose two radically distinct modes of legal reasoning and institutional design. The European framework is characterized by a dense legal lexicon rooted in risk-based typologies, fundamental rights protections, and delegated oversight. In contrast, the American framework privileges executive discretion, innovation enablement, and strategic autonomy, with a lexical field dominated by policy action verbs, institutional actors, and geopolitical registers.

[6] Cf. Ronit Justo-Hanani, 'The politics of Artificial Intelligence regulation and governance reform in the European Union' [2022] Policy Sciences 137 and Neil Guha, Christie M Lawrence, Lindsey A Gailmard, Kit T Rodolfa, Faiz Surani, Rishi Bommasani, Inioluwa Deborah Raji, Mariano-Florentino Cuéllar, Colleen Honigsberg, Percy Liang and Daniel E Ho, 'AI regulation has its own alignment problem: The technical and institutional feasibility of disclosure, registration, licensing, and auditing' [2024] George Washington Law Review 1473.

[7] See, for a more general overview of the different approaches in legislation, among others, Ragnar E. Löfstedt and David Vogel, 'The changing character of regulation: A comparison of Europe and the United States' [2001] Risk Analysis 399; Stanley L. Engerman and Kenneth L. Sokoloff, Institutional and non-institutional explanations of economic differences, in Claude Ménard and Mary M. Shirley (eds.), *Handbook of new institutional economics* (Springer 2005); Bruce E. May, Jeng-Chung V. Chen, and Kuang-Wei Wen, 'The differences of regulatory models and internet regulation in the European Union and the United States' [2004] Information & Communications Technology Law 259.

Despite the presence of a shared conceptual core – terms such as risk, governance, safety, and accountability – the operationalization of AI governance differs substantially across regimes. These divergences may not be merely delegated to textual differences but reveal foundational asymmetries in regulatory rationality, legal authority, and the role of the legislator in technological ordering.

The empirical results presented in this work reveal two fundamentally divergent approaches to legal reasoning and institutional design in the governance of artificial intelligence. The European regulatory model demonstrates a strong orientation towards compliance with formal legal obligations, a structured system of technological oversight, and an overarching emphasis on regulatory predictability and normative standardization. By contrast, the American framework tends to prioritize institutional flexibility, capacity-building, and the strategic enablement of innovation, reflecting a more adaptive and autonomy-driven approach to regulatory design. These distinctions are further substantiated through a PCA of the respective legal texts and policy documents, revealing a high degree of lexical concentration in the European Union's AI Act and a tightly controlled and semantically cohesive regulatory discourse. Conversely, the American corpus displays a more dispersed and heterogeneous lexical distribution, wherein regulatory, economic, and infrastructural themes are interwoven without the same degree of procedural codification or semantic uniformity.

Finally, sentiment analysis uncovers a further axis of divergence in regulatory ethos. The European texts exhibit a tendency toward thoroughness, caution, and risk-aversion, aligning with a rights-centered normative posture. In contrast, the American discourse reflects a more optimistic and expedient tone, privileging innovation and competitiveness. This polarity is particularly pronounced along the axes of speed versus thoroughness, and innovation primacy versus fundamental rights primacy,

thereby illustrating different normative trade-offs across regimes. The study argues that these divergences are not merely textual but symptomatic of foundational asymmetries in regulatory rationality, legal authority, and the role of the state in technological ordering.

The findings underscore the imperative for a more coherent and adaptive regulatory synthesis – one capable of reconciling the normative rigor characteristic of European legal traditions with the procedural agility and market orientation emblematic of the United States' regulatory approach[8]. Such a synthesis does not necessitate the unification or homogenization of distinct legal regimes; rather, it envisions the development of a risk-calibrated, adaptive regulatory paradigm: a governance architecture that modulates normative intensity in proportion to systemic risk, integrates juridical enforceability with mechanisms of innovation governance, and facilitates procedural interoperability without eroding the constitutional pluralism intrinsic to transnational legal orders. In this light, the present inquiry not only delineates existing divergences, but also furnishes both a doctrinal and empirical foundation for the articulation of more dynamic, stratified, and operationally viable models of AI regulation.

From a methodological point of view, the use of algorithmic and computational tools to interrogate legal texts – by identifying structural omissions, disproportionate conceptual weight, and latent semantic asymmetries – suggests their potential as meta-regulatory toolkit that operates alongside, and in support of, risk management

[8] Cf., in this regard, Steven Feldstein, 'Evaluating Europe's push to enact AI regulations: how will this influence global norms?' [2024] Democratization 1049 and Alkandi Ibrahim and Majed A. Helmi 'The impact of strategic agility on organizational performance: the mediating role of market orientation and innovation capabilities in emerging industrial sector' [2024] Cogent Business & Management 1.

strategies and risk-based regulatory systems[9]. These technologies are capable of informing the design, calibration, and critical evaluation of regulatory frameworks themselves. In this light, the comparative analysis developed in this study moves beyond a merely descriptive account of transatlantic divergences and engages with a normative inquiry into the ways in which such tools may influence the governance of emerging technological domains. By examining the legal-linguistic structures embedded in regulatory texts, the study highlights the necessity of a transatlantic regulatory framework that condense the current set of technological recurring patterns, conceptual formulations, and innovation logics that may serve as a resource for the legislator in addressing the evolving demands of AI governance.

Ultimately, this paper refrains from portraying the European or American model as intrinsically superior or characterizing one approach as inherently deficient or dangerously permissive. Rather than adopting a normative stance, the objective is to use computational methods to unpack its explicit legal architectures and the implicit ideological assumptions embedded within each regulatory approach. By combining textual analysis with algorithmic techniques, and integrating legal reasoning with semantic inquiry, this study contributes to the expanding field of computational legal studies, seeking to generate actionable insights for policymakers operating in a context where the traditional pace of legal adaptation is no longer sufficient to meet the challenges posed by rapidly evolving technological systems. In this sense, the project advances not only an empirical investigation of regulatory divergence, but also a methodological reflection on how law can be studied, designed, and reimagined in an age of algorithmic governance.

[9] See Julia Black, 'The emergence of risk-based regulation and the new public risk management in the United Kingdom' [2005] Public law 512.

2. Methodological Framework: Legal Comparison Through NLP

In recent years, the field of comparative legal and economic analysis has witnessed the emergence of new methodological paradigms, driven by the increasing availability of digitized legal corpora and the rapid development of algorithmic tools capable of processing language at scale[10]. While traditional legal scholarship has long emphasized the systematic classification and critical assessment of divergent legal solutions to common normative challenges, a growing body of interdisciplinary research is now exploring how computational techniques can enrich legal analysis[11] by uncovering latent structures, semantic asymmetries, and discursive patterns embedded in legal texts.

This paper adopts a computational approach to the legal analysis of artificial intelligence legislation by applying natural language processing (NLP) techniques to two key regulatory corpora: the European Union's Artificial Intelligence Act (AI Act) and the Executive Orders on AI issued by President Trump. Historically, legal scholarship has focused on identifying the full range of solutions developed by legal systems in response to specific issues, and on evaluating these approaches in a neutral and critically informed manner[12]. Nonetheless, given that recent advances in computational methods provide new opportunities[13] for numerical approaches to

[10] See, in this regard, among others, Zsolt Ződi 'Algorithmic explainability and legal reasoning' [2022] The Theory and Practice of Legislation 67.

[11] Cf., *inter alia*, Jens Frankenreiter and Michael Livermore 'Computational methods in legal analysis' [2020] Annual Review of Law and Social Science 39; Souvik Sengupta 'Legislative Text Analysis from Judicial Case Reports Using Machine Learning' [2024] SN Computer Science 443; Michael Evans, Wayne McIntosh, Jimmy Lin, and Cynthia Cates 'Recounting the courts? Applying automated content analysis to enhance empirical legal research' [2007] Journal of Empirical Legal Studies 1007.

[12] On this point, see Konrad Zweigert and Hein Kötz, *Introduction to comparative law* (Volume I: The Framework, North Holland 1998), 18, in particular.

[13] Computational linguistic tools, in particular, provide innovative means to analyze legal texts at scale, facilitating empirical, data-driven comparisons that extend beyond formal structures to encompass semantic patterns and linguistic emphases in text-intensive regulatory fields such as AI governance.

comparative law[14] we are able to operationalize the comparative inquiry through empirically grounded, language-based metrics able to identify dominant regulatory themes, lexical asymmetries, and the normative weight assigned to specific concepts across jurisdictions, in the attempt of addressing concrete interpersonal, social, and economic problems[15].

The methodology unfolds in three phases: text acquisition and preprocessing; tokenization, lemmatization, and stop-word filtering; vectorization, visualization, and semantic comparison[16]. The source materials consist of the consolidated text of the *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence* – i.e., the AI Act – and the following executive orders: *Removing Barriers to American Leadership in Artificial Intelligence of January 23, 2025; Preventing Woke AI in the Federal Government of July 23, 2025;; Promoting the Export of the American AI Technology Stack of July 23, 2025; Accelerating Federal Permitts of Data Center Infrastructure of July 23, 2025* (Orders)[17]. All documents were stored in plain-text format to ensure compatibility with NLP libraries. The preprocessing stage involved converting all characters to lowercase and removing punctuation and special symbols

[14] See Mathias Siems, *Comparative law* (Cambridge University Press 2022) at p. 202 ff.
[15] This is the specific function of comparative law according to Uwe Kischel, *Comparative law* (Oxford University Press 2019) at p. 8 ff.
[16] For an authoritative analysis of cutting-edge textual analysis methodologies, cf., *ex multis*, Christine P.Chai, 'Comparison of text preprocessing methods' [2023] *Natural Language Engineering* 509; Alexandra George, *Python Text Mining: Perform Text Processing, Word Embedding, Text Classification and Machine Translation* (BPB Publications 2022); Aman Kedia and Mayank Rasu, *Hands-On Python Natural Language Processing: Explore tools and techniques to analyze and process text with a view to building real-world NLP applications* (Packt Publishing Ltd 2020).
[17] These instruments were issued almost concurrently and collectively define the contemporary federal executive framework for AI policy in the United States. Unlike the previous Executive Order 14110, which was normatively dense and administratively layered, these new instruments tends to emphasize economic competitiveness, cultural conservatism, digital infrastructure acceleration, and international technological assertiveness.

using regular expressions, in order to ensure that the analysis focuses exclusively on substantive word content, thus avoiding distortions caused by formatting artifacts[18].

This study uses the *nltk* and *spaCy* Python libraries to tokenize legal texts and segment them into discrete lexical units, to perform lemmatization – a process that normalizes words to their base or dictionary form[19] – and compare semantically equivalent terms exhibiting morphological variation (e.g., uses, using, used). In the context of legal corpora, where terminological precision is closely linked to normative force, lemmatization ensures a more coherent and conceptually unified lexical field for subsequent computational processing[20]. In addition, the analysis incorporated a stop-word filtering phase, targeting functional words and excluding frequently occurring but semantically empty legal verbs and prepositions[21]. Subsequently, this study applied Term Frequency–Inverse Document Frequency (TF-IDF) vectorization method was applied to assess the relative importance of individual terms within each corpus and identify terminological asymmetries that are not merely quantitative but carry interpretive and normative significance[22]. The resulting term-document matrix was visualized via a heatmap and word clouds – generated with Python data visualization libraries *seaborn, wordcloud* and *matplotlib* – both intended to provide immediate, comparative visual inspection of lexical salience across the two corpora.

---

[18] These technical recommendations are set forth in Steven Bird, Ewan Klein, and Edward Loper, *Natural language processing with Python: analyzing text with the natural language toolkit* (O'Reilly Media, Inc. 2009).

[19] For an overview of lemmatization rationale and techniques see Vinay Kumar Pant, Rupak Sharma And Shakti Kundu, 'An overview of stemming and lemmatization techniques' [2024] Advances in Networks Intelligence and Computing, 308.

[20] See, in this regard, Xiang Li, Jiaxun Gao, Diana Inkpen and Wolfgang Alschner, 'Detecting Relevant Differences Between Similar Legal Texts' [2022] Proceedings of the Natural Legal Language Processing Workshop 256.

[21] This procedure is standard, as shown in Souvik Sengupta and Vishwang Dave, 'Predicting Applicable Law Sections from Judicial Case Reports Using Legislative Text Analysis with Machine Learning' [2022] Journal of Computational Social Science 503.

[22] For further discussion of these techniques, see, *inter alia*, Lukáš Havrlant and Vladik Kreinovich, 'A Simple Probabilistic Explanation of Term Frequency–Inverse Document Frequency (TF-IDF) Heuristic (and Variations Motivated by This Explanation)' [2017] International Journal of General Systems 27

To further interrogate the semantic architecture of the two regulatory corpora under study, this paper employs a vector-based representation of lexical content using a *word2Vec* embedding model, followed by dimensionality reduction through *Principal Component Analysis* (PCA) and spatial visualization in three dimensions. This methodological layer seeks to uncover conceptual proximities and divergences that may not be readily visible through frequency-based metrics alone, offering a more nuanced account of the normative semantics embedded in each regulatory framework. The AI Act and the Orders were combined in a unique token stream used to train a *word2vec* model using the *gensim* library[23], positioning words in a continuous vector space where semantic similarity is encoded as spatial proximity. Upon training, the resulting high-dimensional word vectors were then reduced to a three-dimensional coordinate space via *PCA*, allowing for graphical projection and visual interpretation of clusters, outliers, and conceptual alignments between the two regulatory regimes. To differentiate lexical provenance, words were classified into three categories: (i) EU-only tokens, present exclusively in the European corpus; (ii) US-only tokens, unique to the American text; and (iii) common tokens, shared across both. A series of 3D scatterplots was generated to visualize these classes both in isolation and in aggregate.

Al techniques of differential lexical weighting proved to be particularly valuable, as they allowed for a deeper understanding of not only the distribution of emphasis across legal texts, but also the underlying regulatory priorities that such emphasis

[23] As a technical note, the Word2Vec model was trained on a unique token stream using the *Gensim* library, with the following parameter configuration: a vector size of 100, a context window of 5, and a minimum word count of 1. This setup means that each word in the corpus is represented as a 100-dimensional vector, allowing for nuanced semantic relationships to be captured in a dense, continuous vector space. The context window of 5 indicates that the model considers up to five words before and after a target word to learn its contextual associations – an approach that balances local semantic sensitivity with generalization. By setting the minimum word count to 1, the model includes all tokens, even those that occur only once, thus preserving low-frequency but potentially legally significant terms. This is particularly important in legal and regulatory corpora, where rare terms can carry high conceptual weight. This configuration is intended to maximize lexical coverage and preserve the full semantic variability of the legal discourse, which often hinges on subtle terminological distinctions.

reveals. These computational techniques can be insightful for uncovering how legal systems respond to contemporary technological challenges, as they might be able to explain the ways in which different jurisdictions adjust their normative focus to reflect the evolving demands of innovation and the concurrent emergence of new forms of risk. In particular, this approach can reveal how the legal discourse surrounding AIs is evolving in response to new regulatory needs prompted by technological developments, such as the regulation of risk and the legal treatment of issues intrinsically linked to algorithmic decision-making.

In addition to the computational techniques described above, this study incorporated an experimental interpretative layer employing four state-of-the-art large language models: OpenAI's GPT-4.5-turbo (ChatGPT), Microsoft Copilot (Copilot), Google Gemini 2.5 (Gemini), and Mistral Le Chat (Mistral)[24]. The objective is to evaluate the extent to which these systems, when presented with structured and neutral prompts, could generate analytically meaningful observations regarding legal coherence, normative lacunae, and the structural divergences between the two regulatory regimes across the Atlantic. This experimental layer was intended to complement the doctrinal analysis by assessing the interpretative capacity of LLMs when applied to complex legal corpora. The outputs were then evaluated for analytical depth, internal consistency, responsiveness to the prompts, and the degree to which the responses mirrored or deviated from conventional legal reasoning.

To ensure methodological consistency and replicability, each large language model was provided with the complete text of both regulatory instruments under

[24] The transcripts of interactions with the various AI models are accessible via the following links: Mistral, https://chat.mistral.ai/chat/79fb9892-f182-464b-98c1-c80b50142d40; ChatGPT, https://chatgpt.com/share/689331d8-6c28-8004-9397-d5ee73994939; Gemini, https://g.co/gemini/share/557e9967fe61; Co-pilot, https://copilot.microsoft.com/shares/JaacKL15e8RxjckzvZra6.

examination and tasked with responding to a standardized, multi-part prompt designed to elicit a granular document-level assessment focused on three dimensions: legal strengths (e.g., clarity, enforceability, conceptual robustness); weaknesses (e.g., ambiguity, interpretative difficulty, implementation burdens); and regulatory gaps (i.e., areas insufficiently addressed or omitted entirely). The prompts were the following:

> 1. "Please analyze the following legal text in depth. Go beyond technical identification of strengths and weaknesses: provide a critical and evaluative assessment, offering well-reasoned value judgments on the text's overall quality, coherence, and legal soundness. Identify any strengths, weaknesses, or gaps in legal reasoning, structure, or support, and explain why they matter. Pay attention to clarity, consistency with applicable law, completeness, and argumentative strength. Highlight any ambiguous language or provisions that may hinder interpretation or implementation. Also, point out any missing provisions or underdeveloped areas that leave important issues unresolved or undefined. Where appropriate, suggest substantive improvements or alternatives."
>
> 2. "Now that you have analyzed each legal text individually, provide a detailed and critical comparative analysis addressing the following:
> 1. Comparison of Strengths: Clearly identify the unique strengths of each text. Evaluate whether the strengths of one document complement, reinforce, or compensate for the weaknesses of the other. Assess the overall legal and practical effectiveness of each document in its context.
> 2. Comparison of Weaknesses: Examine whether the weaknesses in one text are also present in the other, or whether one document provides better solutions or safeguards. Comment on the significance of these weaknesses in terms of legal clarity, enforceability, or doctrinal coherence.
> 3. Comparison of Gaps and Omissions: Identify overlapping or distinct gaps in both documents. Where possible, determine whether one text addresses what the other fails to cover, and discuss whether such omissions impact legal certainty or regulatory completeness.
> 4. Contradictions and Tensions: Detect any contradictions, inconsistencies, or normative tensions between the two texts. Analyze their implications for legal interpretation, application, or integration.
> 5. Synthesis and Harmonization: Provide a thoughtful synthesis of how the two texts could be reconciled or harmonized. Offer concrete suggestions for improving consistency, filling gaps, or aligning provisions to enhance legal coherence and effectiveness.
> Where relevant, support your analysis with legal reasoning, examples, or references to legal principles, and maintain a critical and evaluative tone throughout."

Preliminary findings suggest that, while LLMs vary significantly in their interpretative accuracy and rhetorical precision, they are capable – under carefully controlled conditions – of identifying latent regulatory tensions, inconsistencies in scope or terminology, and even implicit normative hierarchies embedded in legal texts. This AI-assisted layer of interpretation added an additional epistemic dimension to the analysis. Given that, when appropriately prompted, indeed, large language models prove to be capable jurilinguistic tools[25], proficient in identifying latent regulatory asymmetries, conceptual discontinuities, and prospective avenues for harmonization and legal clarity. In this regard, LLM can serve as a powerful methodological tool for legal scholarship operating in computational legal studies, seeking to apply quantitative, algorithmic methods to traditionally qualitative domains of legal reasoning[26]. The results support a triangulated perspective on the underlying legal structures and normative assumptions inherent in each legal regime's approach to AI, offering not only an additional epistemic dimension but also new interpretive heuristics and innovative regulatory solutions in the governance of technology[27].

Furthermore, this study incorporates a quantitative sentiment analysis layer as an auxiliary methodology designed to assess evaluative tone in AI-generated commentary on both regulatory frameworks. This supplementary quantitative analysis provides a statistical proxy for detecting how different LLMs frame the regulatory text in terms of positive or negative evaluative language. To this end, the full-text outputs of the four LLMs interviews were subjected to automated sentiment scoring using the *Valence*

[25] For an overview of the methodological approaches followed in the analysis, see Stanisław Goźdź-Roszkowski, 'Corpus Linguistics, Methodology of Jurilinguistics' in Anne Wagner and Aleksandra Matulewska (eds), *Research Handbook on Jurilinguistics* (Edward Elgar Publishing 2023).
[26] On this topic see, among others, Ryan Whalen, *Computational Legal Studies: The Promise and Challenge of Data-Driven Research* (Edward Elgar Publishing 2020)
[27] Cf., on this matter, Gerd Gigerenzer and Christoph Engel (eds), *Heuristics and the Law* (MIT Press 2006) and Michael Guihot, Anne F. Matthew and Nicolas P. Suzor, 'Nudging Robots: Innovative Solutions to Regulate Artificial Intelligence' (2017) Vanderbilt Journal of Entertainment and Technology Law 385

Aware Dictionary and *sEntiment Reasoner* (also known as *VADER*) algorithms, a lexicon-based and rule-based tool specifically calibrated for analyzing sentiment in English textual data. *VADER* computes polarity scores across three primary dimensions (positive, negative, neutral) and a score that reflects overall sentiment. For the purposes of this analysis, only the positive and negative scores were retained, and were subsequently normalized by computing the ratio of each sentiment class to the total polarity load (i.e., positive + negative). This normalization controls for overall sentiment density and allows for comparative evaluation across models regardless of text length or verbosity. Each LLM's output was processed as a single document, and the resulting normalized scores were then visualized in a comparative bar chart.

3. Quantitative results: Lexical Divergences and Regulatory Emphases

The results of the computational analysis reveal substantively significant and structurally profound differences in the semantic architecture of the two regulatory corpora under examination. While both the AI Act and the Orders on AI ostensibly aim to address a common regulatory object – namely, the governance, regulation, and oversight of AI systems – the respective texts of these legal instruments exhibit notable divergences in their structural composition, lexicon, and prioritization of legal concepts. These variations point to fundamentally distinct regulatory grammars, each reflecting divergent normative frameworks, institutional presumptions, and policy orientations that inform their underlying legal philosophies. The word cloud visualizations – see Figure 1 – offer an initial yet highly informative quantitative representation of the dominant terminologies within both legal texts, highlighting the most frequently employed terms, thereby facilitating an immediate comparative analysis of the linguistic emphasis placed by each framework on particular concepts.

Figure 1

Note: The figures represent word-clouds generated from the full texts of the EU Artificial Intelligence Act and the US Executive Orders on AI, based on a lemmatized, stopword-filtered corpus processed using *spaCy* and *NLTK* Python libraries. Tokens were normalized to their lemma forms, and high-frequency semantically vacuous terms were excluded to improve conceptual salience. The resulting lexical distributions were visualized using the *WordCloud* and *Matplotlib* libraries, offering a comparative overview of the dominant regulatory lexicon employed in each legal instrument.

The word cloud associated with the EU AI Act reveals a dense aggregation of legally operational terminology, which serves as a clear reflection of the Act's highly structured and prescriptive regulatory framework. Dominant terms, such as "system," "high-risk," "provider," "authority," "compliance," "purpose," "conformity," and "person," *inter alia*, underscore the Act's foundational logic of *ex ante* control, tiered

obligations based on risk assessment, and the harmonization of legal subjectivity. The prominence of actor-specific designations[28] (e.g., "provider," "deployers"), regulatory qualifiers (e.g., "high-risk," "general purpose"), and technical and procedural mandates (e.g., "conformity assessment," "surveillance") captures the allocation of responsibilities and enforceable duties within a risk-based horizontal regulatory paradigm, which is fundamentally designed to safeguard fundamental rights[29]. This regulatory model, organized around distinct and predefined risk categories, aligns with the broader tradition of product safety legislation, wherein regulatory mechanisms are structured to prevent harm before it materializes, reflecting a precautionary approach to risk mitigation, according to predetermined risk tiers[30].

In contrast, the lexical distribution within the US regulatory corpus is predominantly characterized by terms such as "presidential," "order," "policy," "national," "security," "America," and "international," thereby indicating a distinctly executive-driven and sovereignty-oriented approach to the governance of artificial intelligence. Rather than structuring its terminology around concepts of risk, legality, or enforceability, the text foregrounds institutional referents (e.g., "president," "agencies," "departments"), geopolitical markers (e.g., "international," "foreign," "America," "allies"), and thematic imperatives related to technological leadership and national interest (e.g., "export," "industry," "infrastructure," "deployment"). The predominance of action-oriented and directive verbs—such as "promote," "support," "accelerate," and "ensure"—further underscores the document's primary role as an

[28] On this specific topic, see Mitja Sienknecht and Antje Vetterlein, 'Conceptualizing Responsibility in World Politics' [2024] *International Theory* 26.
[29] Cf. Martin Ebers, 'Truly Risk-Based Regulation of Artificial Intelligence: How to Implement the EU's AI Act' [2024] European Journal of Risk Regulation 1; Michael Veale and Frederik Zuiderveen Borgesius, 'Demystifying the Draft EU Artificial Intelligence Act: Analysing the Good, the Bad, and the Unclear Elements of the Proposed Approach' [2021] Computer Law Review International 97.
[30] This similarity is discussed in Marco Almada and Nicolas Petit 'The EU AI Act: Between the rock of product safety and the hard place of fundamental rights' [2025} *Common Market Law Review* 85.

agenda-setting instrument, rather than a prescriptive legal text. This aligns with the broader American regulatory philosophy, which tends to favor flexible frameworks and general guidelines over rigid legal mandates[31], thus minimizing the imposition of excessive regulatory constraints on innovation.

Notably absent from the lexical field of the US corpus are terms indicative of binding obligations or structured compliance regimes applicable to private sector actors[32]. Instead, the discursive emphasis on "national security" reflects a model of governance that centers on the strategic positioning of the state within the global technological market, with a particular focus on safeguarding national interests in the military and security sectors[33]. The executive-centric approach is intended to favor innovation-oriented coordination strategies, public-private partnerships, and directives grounded in geopolitical priorities[34]. The US regulatory architecture is thus legally thin but institutionally dense: it establishes authority not through the formal articulation of rights, duties, or enforceable obligations, but through vertical delegation, administrative hierarchy, and strategic directionality, thereby reflecting a managerial, and at times instrumental, conception of AI governance – one in which state power is exercised less through constraint and prohibition, and more through centralization, policy signaling, and executive fiat. This model aligns with a broader tradition within

[31] See Ciro Sbailò, 'Governing Artificial Intelligence: Technological Leadership and Regulatory Challenges in an Era of Exponential Growth' [2024] DPCE Online 275.
[32] This observation is hardly novel, as it has been previously noted by Colin Scott, 'Private regulation of the public sector: a neglected facet of contemporary governance' [2002] Journal of Law and Society 56.
[33] From this perspective, see the contribution of Daniel J. Mallinson, Icon,Lauren Azevedo,Eric Best, and Pedro Robles, 'Artificial intelligence policy, the Trump Administration, and federalism' [2025] Administrative Theory & Praxis 1.
[34] It has been argued, however, that the approach to AI and its emphasis on competitiveness scarifies trust and safety, fostering "chaos and uncertainty", undermining critical elements of ethical governance. See, in this regard, Susan A. Aaronson and Michael Moreno 'Regrets of the Tech Bros: In a land ruled by the law of the jungle' [2025] The International Economy 39.

the American administrative state, which tends to prioritize technological innovation, sectoral flexibility, and post hoc remedial action over *ex ante* legal formalism[35].

The lexical divergences between the US and EU corpora thus reveal two distinct regulatory philosophies: while the EU text leans toward legal formalism, normative closure, and the imposition of enforceable legal obligations rooted in rights-based discourse, the US document reflects a managerial approach to governance, primarily situated within the domain of soft law and administrative discretion where legal enforceability remains limited[36] and the primary focus is placed on institutional adaptability, strategic coordination, and responsiveness to emerging challenges.

A weighted and more systematic analysis of lexical salience across documents is reported in the TF-IDF heatmap (see Figure 2), which visually represents the prominence of key terms within each regulatory corpus.

Figure 2

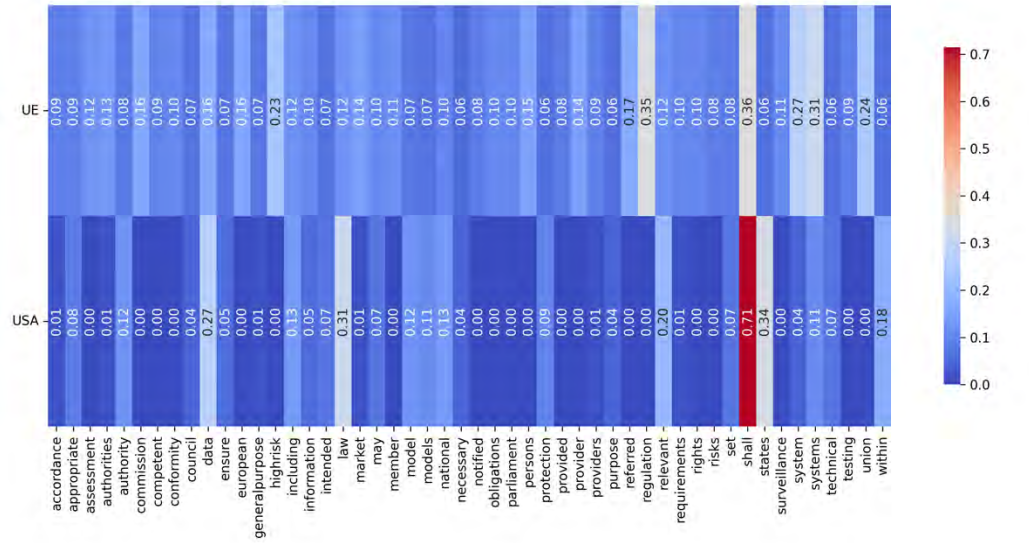

Note: The figure shows a Term Frequency–Inverse Document Frequency (TF-IDF) heatmap comparing the EU Artificial Intelligence Act and the US Executive Order on AI. The matrix was generated using scikit-learn's TfidfVectorizer on a preprocessed, lemmatized, and stopword-filtered corpus of the two regulatory texts. The heatmap – created using the *seaborn* library in Python – displays the relative importance of the fifty most distinctive terms in each corpus, illustrating divergent patterns of lexical prioritization and normative emphasis across the two jurisdictions.

[35] In this regard, see Julia Black, Martin Lodge and Mark Thatcher (eds), Regulatory Innovation: A Comparative Analysis (Edward Elgar Publishing 2006).

[36] In support of this view, see Cary Coglianese, 'People and Processes: AI Governance Under Executive Order 14,110' [2023] *Administrative and Regulatory Law News* 9

The TF-IDF heatmap visually represents the distribution of lexical salience across the two normative corpora, highlighting the relative prominence of specific terms within the AI Act and the Orders. As a computational technique, TF-IDF is particularly well-suited for juridical comparison, as it assigns statistical weight to terms that are conceptually central within a given corpus while remaining relatively rare across the broader comparative dataset. This method enables the identification of regulatory asymmetries not only in the frequency of terminological usage but also in the deeper structural and normative dimensions of the legal frameworks. By focusing on the statistical significance of terms, TF-IDF uncovers insights into the differing legal priorities and underlying regulatory logics that shape each corpus.

In the EU row, a distinct cluster of high-weighted terms – such as "system," "high-risk," "requirements," "provider," "conformity," "compliance," "surveillance," "assessment," and "data" – clearly delineates – once again – a regulatory framework centered on formal legal obligations. These terms are not arbitrarily distributed; rather, they are semantically and normatively interdependent, collectively forming the foundational elements of a risk-based regime that imposes binding duties on defined actors operating within specified categories of technological functionality[37]. The prominence of the term "high-risk," in particular, underscores the EU's specific focus on addressing the most hazardous AI systems, reflecting a regulatory taxonomy that prioritizes AI technologies with the potential for significant harm. This regulatory framework is rooted in the broader legal culture of the EU, which emphasizes the categorization and mitigation of risks, with a particular emphasis on technologies deemed to pose serious threats to public safety, fundamental rights, or security.

[37] There is a concern, on this topic, on regulatory gaps and potential overlaps, as noted in Custers Bart Custers, Henning Lahmann, and Benjamyn I. Scott, 'From liability gaps to liability overlaps: shared responsibilities and fiduciary duties in AI and other complex technologies' [2025] AI & society 1.

The prominence of terms such as "surveillance," "conformity," and "compliance" further confirms a normative infrastructure that hinges on enforceable procedural duties and delegated technical oversight, consistent with a legal culture that privileges predictability, standardization, and administrative enforceability[38]. Together, these terms point to a model of legality that operationalizes accountability primarily through administrative mechanisms – e.g., certifications, technical documentation, and oversight by notified supervisory bodies and regulatory authorities – rather than relying on discretionary standards or voluntary compliance. This regulatory framework, with its emphasis on procedural rigor, seeks to ensure that AI systems meet established criteria before they are deployed, thereby safeguarding against potential risks and reinforcing public trust in AI technologies[39].

By contrast, the lexical profile of the US corpus reveals a markedly distinct regulatory approach, structurally divergent from that of the EU. Dominated by terms such as *"national," "international," "development," "promote,"* and *"support,"* the US text is organized more around policy direction than juridical enforceability. These terms reflect a regulatory framework that places a premium on the role of the federal government as a facilitator of technological innovation, through strategic leadership in the global race on AI[40], rather than through the imposition of formal, codified legal rules. Indeed, the recurring prominence of terms like *"promote," "support,"* and *"development"* signals a deliberate regulatory posture aimed at encouraging technological growth, industrial expansion, and international competitiveness. Rather

[38] See, in this regard, Sorina Serban-Barbu, 'Public Administration Regulations Codification in the European Union' [2014] Comparative Law Review 81.
[39] This is, indeed, the scope and objective of the AI Act, according to Johann Laux, Sandra Wachter, and Brent Mittelstadt, 'Trustworthy artificial intelligence and the European Union AI act: On the conflation of trustworthiness and acceptability of risk' [2024] Regulation & Governance 3.
[40] See, in this perspective, Stefka Schmid, Daniel Lambach, Carlo Diehl and Christian Reuter, 'Arms race or innovation race? Geopolitical AI development' [2025] Geopolitics 1.

than seeking to constrain innovation through rigid legal duties or exhaustive regulatory oversight, the US framework is oriented toward creating an environment conducive to innovation and close collaboration between big techs and the administration[41].

From a technical point of view, salient lexical asymmetry emerges from the prominence of the modal verb shall in the US corpus, which registers the highest TF-IDF score in the entire heatmap (0.71), while occupying a notably more marginal position in the EU corpus (0.36). At first glance, this distribution may seem counterintuitive, given the prescriptive nature of the EU text. However, a closer analysis reveals that "shall" in the US corpus primarily functions as a formalistic expression of presidential directive authority, embedded in executive mandates and policy imperatives[42]. In contrast, the EU corpus – despite its strong regulatory intensity – typically encodes obligations through institutionally distributed norms, technical requirements, and definitional clauses, often avoiding modal expressions[43] in favor of direct legal provisions, embedded in structured articles and cross-referenced mechanisms, rather than articulated through high-frequency imperatives.

Taken together, the heatmap not only visualizes patterns of lexical differentiation but also reveals a deeper regulatory bifurcation: on the one hand, a European model oriented toward a rights-protective legal structure, characterized by codified obligations and administrative enforceability; on the other, a US model grounded in political directionality, strategic promotion, and institutional discretion[44]. The TF-IDF

[41] See, on this topic, the comments enclosed in Jeff Tollefson Witze and Max Kozlov, 'Science under Trump 2.0: The likely winners and losers' [2025] Nature 533.
[42] In the context of executive discretion, elements supporting this view are explored in detail by Zachary S. Price, 'Enforcement Discretion and Executive Duty' [2014] Vanderbilt Law Review 671.
[43] See, in this regard, Giuliana Garzone, 'Variation in the Use of Modality in Legislative Texts: Focus on Shall' [2013] Journal of Pragmatics 68.
[44] It is worth noting that the previous US administration's Executive Order – since repealed by the current legislature – differed structurally in several respects, yet articulated commitments to algorithmic accountability, bias mitigation, and pre-deployment stress testing (commonly referred to as "red-teaming") in a manner broadly aligned with the European Union's AI Act. Significantly, both regulatory

analysis thus substantiates the broader thesis of this work: that divergent conceptions of legality, risk, and public authority are not merely abstract theoretical differences, but are materially encoded in the semantic architecture of contemporary AI regulation.

Finally, to complement the previously discussed semantic similarity analysis, this study employed a vector-based representation of lexical content through high-dimensional word embeddings, using the Python Word2Vec library to generate a graphical projection and spatial visualization (see Figure 3) of semantic relationships and conceptual distances between terms.

Figure 3

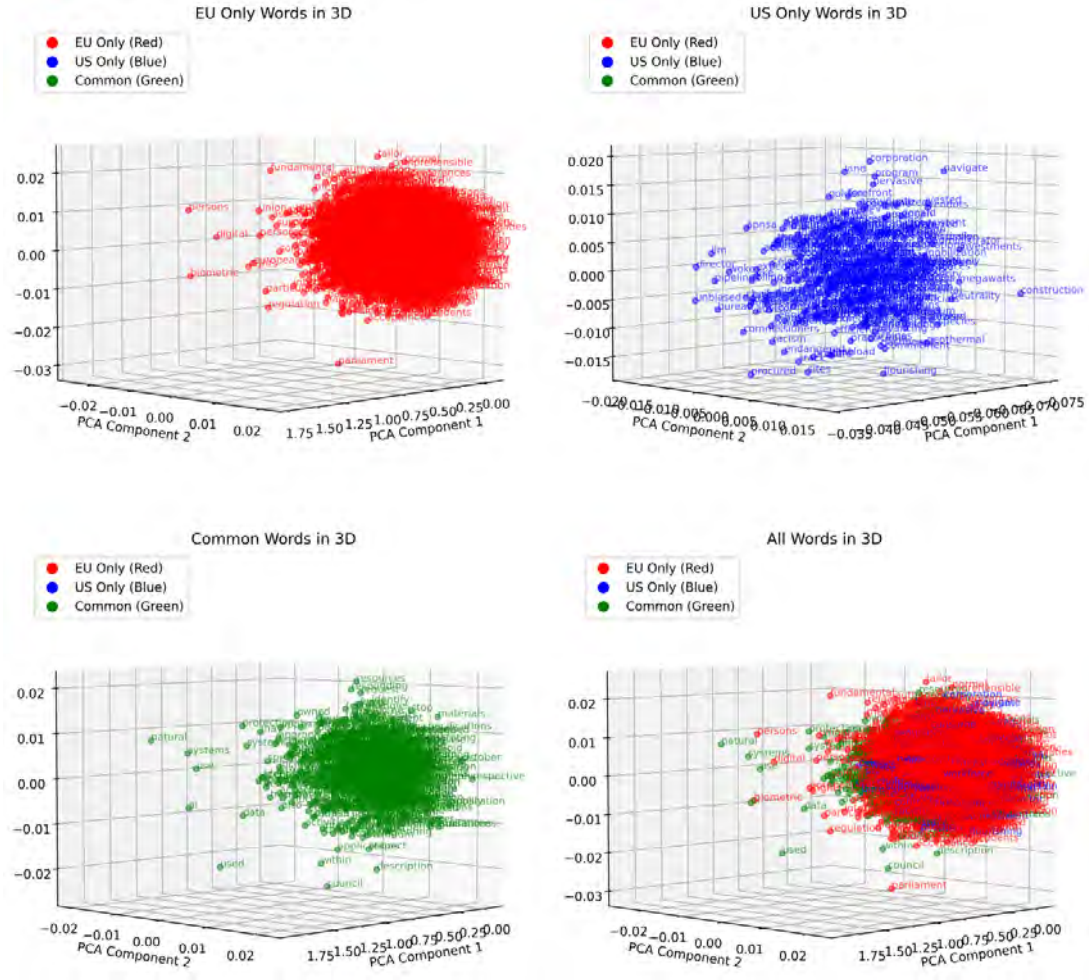

instruments incorporate carve-outs for national security applications – a concession that has elicited criticism from civil liberties organizations. See Trevor Odelberg, 'Understanding the Future of Artificial Intelligence Governance: Comparing the EU AI Act and US Executive Order on Safe AI' [2024] Gerald R. Ford School of Public Policy Research Report 1.

Note: The figure shows a 3D lexical embedding of regulatory vocabularies from the EU AI Act and US Executive Order(s) on AI, generated via Word2Vec (vector size = 100) and reduced to three principal components using PCA. Tokens are colored by jurisdictional origin: red (EU-only), blue (US-only), green (shared). The visualization highlights semantic clustering, lexical asymmetries, and conceptual divergence between the two regulatory systems.

The vectorized semantic mapping produced via Word2Vec embedding and PCA-based dimensionality reduction illustrates with notable clarity the structural divergences between the European and American regulatory lexicons on AI. The EU-exclusive cluster (red, top-left panel) emerges as a dense and spatially compact aggregation, centered around a lexicon that is both juridically formalized and institutionally anchored. Prominent terms such as "fundamental", "persons", "digital", "biometric", "parliament", "comprehensible", and "authorization" delineate a governance discourse rooted in rights-based constitutionalism, procedural transparency, and explicit legislative authority. The presence of "fundamental" and "persons" reflects the embedding of AI regulation within the broader framework of the EU Charter of Fundamental Rights and its associated jurisprudence[45], where the human subject remains the primary normative referent.

The lexeme "biometric" links the regulatory text to the domain of advanced data-driven technologies, yet its position within the EU-exclusive cluster suggests that it is not employed as a neutral technical descriptor but it forms a legally defined category subject to stringent European-only procedural safeguards[46]. The inclusion of "parliament" and "authorization" further highlights the EU model's reliance on democratic legitimacy and delegated legal authority within a complex, multi-level

[45] Cf., in this regard, Francesco P. Levantino and Federica Paolucci, 'Advancing the Protection of Fundamental Rights Through AI Regulation: How the EU and the Council of Europe are Shaping the Future' [2024] European Yearbook on Human Rights 3, and Isabel Kusche, 'Possible harms of artificial intelligence and the EU AI act: fundamental rights and risk' [2024] Journal of Risk Research 1.

[46] On this specific matter, see Bilgesu Sumer, Natalia Menéndez González, Abdullah Elbi, Catherine Jasserand, Jan Czarnocki and Els J. Kindt, 'AI Acts' Ripple Effect on Biometric Data: Harmonising or Fragmenting the Regulation of Biometric Data' in Kostina Prifti, Esra Demir, Julia Krämer, Klaus Heine, Evert Stamhuis (eds.) *Digital Governance: Confronting the Challenges Posed by Artificial Intelligence* (165-181, TMC Asser Press | 2024).

governance framework. Meanwhile, the salience of "comprehensible" signals an explicit regulatory commitment to intelligibility and transparency – not merely as aspirational values, but as enforceable obligations embedded within a constitutional architecture that demands uniform pluralism and administrative coherence[47].

The US-exclusive cluster (blue, top-right panel) appears markedly more dispersed, reflecting a heterogeneous lexical field in which regulatory, economic, and infrastructural themes coexist without the same degree of formal integration. This fragmentation mirrors the procedural informality and sectoral pluralism that characterizes the American regulatory state[48]. Central terms include "corporation", "program", "construction", "pervasive", "finalized", "procured", "unbiased", "director", and "bureau". The prominence of "corporation" and "program" underscores the centrality of private-sector actors and programmatic implementation mechanisms in the American model, where governance often advances through contractual arrangements, public–private partnerships, and executive-led initiatives rather than through codified legislative mandates.

The terms "construction", "finalized", and "procured" point to a regulatory discourse embedded in project delivery, procurement procedures, and infrastructural expansion – elements that align with a policy orientation favoring capacity-building, implementation speed, and strategic rollout over *ex ante* legal control. The inclusion of "unbiased" serves as a normative marker, yet its semantic positioning within the cluster suggests an evaluative, performance-based understanding of fairness, rather than a rights-based enforcement logic. Finally, the co-occurrence of "director" and "bureau" reflects an executive-branch-centric nature of the US model, in which

[47] See, in this regard, Andrew Williams, The Ethos of Europe: Values, Law and Justice in the EU (Cambridge University Press 2010).
[48] To this effect, see David Colander and Harry Landreth, 'Pluralism, Formalism and American Economics' in Edward Fullbrook (ed), *Pluralist Economics* (26 ff., Edward Elgar Publishing 2008).

administrative agencies and appointed officials are entrusted with wide discretion in operationalizing regulatory priorities and shaping sectoral governance norms.

The common vocabulary cluster (green, bottom-left panel) occupies an intermediary semantic space between the two regulatory regimes, comprising terms such as "natural", "used", "council", "description", "AI", "resources", "stop", "transparency", and "public". At first glance, this set of terms may suggest the emergence of a shared transatlantic lexicon. However, their spatial positioning within the embedding model – and, crucially, their semantic proximities (see Figure 3) – reveal divergent functional roles and normative weights across the two corpora.

In the EU corpus, terms such as "transparency" and "public" are normatively charged, embedded within a framework of enforceable obligations related to disclosure, administrative oversight, and participatory governance. By contrast, in the US corpus, these same terms are more frequently situated in programmatic or communicative contexts, functioning as aspirational policy objectives rather than legally binding duties.

The term "council" likewise illustrates this divergence: within the EU context, it typically denotes formal institutional bodies endowed with defined legal competences, whereas in the US regulatory discourse it more often refers to advisory groups or inter-agency coordination mechanisms, lacking direct normative authority. Even ostensibly neutral terms such as "resources" and "description" acquire distinct operational valences—linked in the EU to capacity-building obligations and compliance infrastructure, and in the US to the allocation of administrative inputs within executive-led programs.

Before turning to the qualitative evaluations elicited from the language models themselves, this section introduces a preliminary computational sentiment analysis,

designed to map the evaluative polarity assigned by each large language model (LLM) to the two regulatory corpora under examination. As illustrated in Figure 4, this sentiment analysis functions not as an isolated metric, but as a preparatory heuristic tool able to contextualize the models' subsequent discursive outputs by identifying underlying polarity patterns that may influence their interpretive responses.

Figure 4

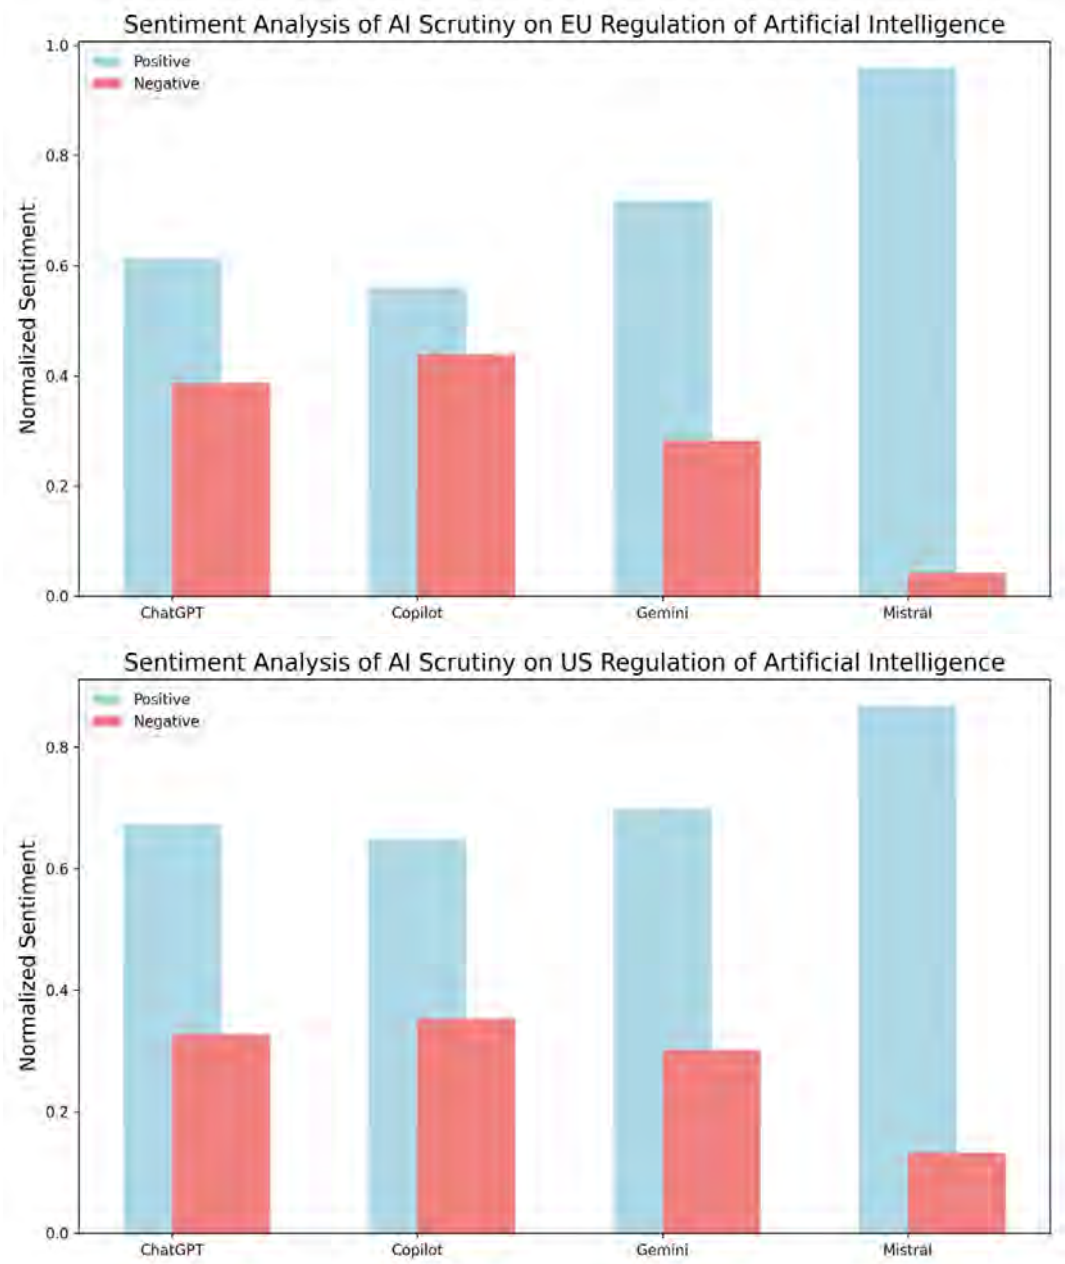

Note: The figures display a sentiment analysis pertaining to the EU AI Act and the US Executive Orders, conducted using four Large Language Models (LLMs): ChatGPT, Copilot, Gemini, and Mistral. The charts visualize the normalized sentiment polarity (positive vs. negative) expressed by each LLM towards the two regulatory frameworks. This quantitative sentiment analysis provides an empirical basis

for interpreting the LLMs' perceptions and evaluations regarding the characteristics, strengths, and critical areas of the examined regulations.

The results indicate a clear prevalence of positive sentiment (light blue) over negative sentiment (light red) across both legal instruments. In all four models, the proportion of positive sentiment consistently exceeds that of negative sentiment, suggesting that—within the computational parameters of this analysis—favorable evaluations outweigh unfavorable ones in both cases. Yet, beyond this general positivity bias, the sentiment analysis reveals persistent structural divergences in how the EU and US regulatory frameworks are perceived by each model. Mistral, for instance, expresses a markedly higher degree of positive sentiment toward the AI Act, while Copilot, ChatGPT, and Gemini tend to allocate greater positive weight to the Orders.

Nonetheless, the differences in positivity allocation highlight that the LLMs do not treat the two frameworks as equivalent or interchangeable. Mistral's comparatively stronger positive response to the AI Act may indicate an appreciation for its detailed legal formalism, structured risk-based approach, or rights-centered governance ethos. Conversely, the higher positive sentiment toward the Orders by Copilot, ChatGPT, and Gemini might reflect an alignment with or a recognition of the US model's managerial, innovation-friendly stance, emphasizing flexibility, executive leadership, and sectoral pragmatism. These divergent sentiment profiles underscore how underlying regulatory philosophies – i.e., formal legalism versus managerial discretion – might be perceptible even in computational sentiment outputs[49]. They also prefigure how

[49] It is worth acknowledging the potential for a location bias influencing the sentiment results, as Mistral is developed by a European-based entity, whereas Copilot, ChatGPT, and Gemini are products of US-based organizations. This geographic distinction may partially explain the divergent sentiment patterns observed. However, given the limited sample size and the complexity of LLM training data and architectures, there are insufficient grounds to draw definitive causal conclusions or broad generalizations regarding the impact of developer location on model evaluation biases.

individual LLMs might prioritize or critique different facets of AI governance when prompted in the subsequent qualitative interviews.

To complement the quantitative sentiment analysis, a qualitative examination of the large language models' responses was conducted. This approach aims to provide deeper insights into how each model interprets and evaluates the AI Act and the Orders on beyond mere polarity scores. This qualitative layer enriches the overall analysis by revealing the subtleties and complexities that quantitative measures alone may overlook. In all interviews, all LLM converge in their identification of systemic weaknesses in all interaction.

The qualitative interviews, with the four LLMs, show distinct hermeneutic readings of the AI Act and of the Orders, rendering separated interpretations and highlighting different dimensions of regulatory intelligibility[50].

ChatGPT, for instance, treats the European AI Act as an "attempt to comprehensively regulate AI within its territory" grounded on a "risk-based regulatory framework, with graded obligations depending on the level of risk AI systems pose", that is "carefully grounded in EU Treaties" and it is "consistent with the Charter of Fundamental Rights". In ChatGPT's view, the AI Act is "impressively detailed and systematic" and "avoids the "patchwork regulation" problem seen in some national initiatives" and offers a strong protection of fundamental rights, being "built on human-centric values". The "risk-calibrated approach avoids overregulation while still imposing serious controls where harms are most likely". Nevertheless, the biggest AI Act weakness is the "lack of technical thresholds**"** – which could lead to interpretive inconsistencies – and the uncertainty surrounding its implementation with the serious risk of "becoming a symbolic framework rather than a practically effective one"

[50] Please refer to note 24 for the various hyperlinks corresponding to the distinct conversations conducted with each language model.

because of a “too broad high-risk classification” and “limited focus on accountability mechanisms”.

Regarding the US Orders, ChatGPT frames them as coherent instruments of executive policymaking – framing AI governance as a matter of “coordinated federal action” aimed at fostering institutional collaboration and administrative feasibility – they nonetheless exhibit significant shortcomings. Chief among these is a marked lack of specificity, representing a “missed opportunity to articulate substantive guidance on the ethical development of AI, technical safety standards, or technical performance benchmarks”. Moreover, the Orders are “unusually polemical” in tone for legal instruments of their kind, at times adopting politically charged language that risks interpretive ambiguity and heightened exposure to litigation. Particularly concerning is the apparent framing of certain social science theories as “inherently false or harmful,” a move that imports vague, normatively laden assertions into a legal text, thereby increasing the risk of confusion, selective enforcement, or politicized misuse during implementation. Equally notable is the absence of any reference to sustainability considerations, emissions reduction targets, or energy efficiency standards applicable to large-scale AI infrastructure, such as data centers. This omission suggests a deregulatory posture shaped more by ideological considerations than by a commitment to technical rigor or public accountability.

Copilot’s analysis adopts an administrative-law perspective, scrutinizing both instruments with particular attention to their drafting architecture, implementation mechanisms, and embedded accountability structures. Within this analytical framework, the AI Act is construed as a comprehensive statutory regime that “ambitiously encompasses every facet of AI”, comprising “more than 170 recitals and 113 articles”, supplemented by extensive annexes. The regulation introduces a risk-

based approach, calibrating legal obligations in accordance with the potential harm posed by different AI systems. However, the Act's breadth gives rise to structural concerns. It attempts to integrate numerous sector-specific amendments, anticipatory recitals addressing a vast array of potential use cases, and simultaneous regulation of so-called "general-purpose" AI models. This legislative layering risks engendering regulatory complexity to the point of paralysis. Moreover, Copilot notes that the Act grants market surveillance authorities sweeping powers of remote access to source code, datasets, and system logs. While framed as a means of ensuring compliance, such provisions raise legitimate concerns regarding the protection of proprietary information, the chilling of investment, and the erosion of trade secret safeguards. Finally, the harmonized penalty structure – providing for fines of up to 7% of global annual turnover – is applied uniformly across sectors. The regulation affords only limited mitigation for smaller entities, with the SME exemption functioning more as a symbolic gesture than a meaningful accommodation.

Copilot's analysis approaches both instruments through an administrative-law lens, paying close attention to drafting architecture, implementation pipelines, and accountability by design. In its reading, the AI Act emerges as a comprehensive statutory framework that "ambitiously covers every facet of AI", with "over 170 recitals and 113 articles (plus annexes)", providing tailored obligations according to a risk-based approach. The regulation, however, "attempts to layer in dozens of sector-specific amendments", recitals on every imaginable use, and simultaneous regulation of "general-purpose" models. This risks paralysis. Furthermore, market surveillance authorities get expansive remote-access rights to source code, data and logs, potentially chilling investment and revealing trade secrets. Finally, Harmonization of

fines (up to 7 % turnover) applies across all sectors, with little differentiation for smaller players beyond a token SME discount.

Copilot's examination of the U.S. Executive Orders is equally rigorous, emphasizing their operational reliance on interagency delegation, prescriptive implementation timetables, and federal procurement as a lever of regulatory influence. At the same time, Copilot identifies a range of legal and practical challenges arising from imprecise or ideologically loaded terminology. Notably, the "labeling of diversity, equity, and inclusion (DEI) as an 'ideology' and equating it with 'suppression of factual information' mischaracterizes well-established practices" and misrepresent established anti-discrimination frameworks. Such framing, Copilot argues, risks "chilling legitimate bias-mitigation efforts" and undermining the development of inclusive AI governance. These ambiguities raise broader concerns regarding administrative consistency, judicial reviewability, and the meaningful inclusion of stakeholders in implementation processes. Furthermore, the Orders are noted to exhibit "minimal civil-rights oversight" and lack a coherent procedural framework to coordinate agency action, thereby exacerbating the potential for interagency fragmentation and interpretive divergence. Finally, Copilot observes that the current U.S. approach is marked by the absence of enforceable civil liability mechanisms, fails to mandate third-party audits or certification for bias-mitigation measures, and provides only limited avenues for public consultation. Collectively, these omissions weaken transparency, accountability, and democratic legitimacy of American AI regulation.

Gemini characterizes the AI Act as a "comprehensive, principled, and human-centric approach to AI governance." Echoing the position noted by ChatGPT, Gemini underscores that the regulation is explicitly grounded in the premise that AI should

function as a “human-centric technology,” developed in accordance with Union values, fundamental rights, and the Charter of Fundamental Rights of the European Union. A notable strength, according to Gemini, lies in the Act’s commitment to legal coherence: the regulation expressly positions itself as “complementary” to existing EU legal frameworks governing data protection, consumer protection, and product safety. This integrative approach is seen as reinforcing consistency across the broader corpus of Union law. Uniquely among the models surveyed, Gemini highlights the recitals as providing a “robust set of definitions for key terms,” which it views as critical to ensuring legal certainty. However, Gemini also acknowledges several limitations. Chief among them is the potentially “ambiguous language used in key prohibitions”, in particular in delineating "unacceptable" AI practices. In addition, the regulation is critiqued for its underdeveloped safeguards surrounding exemptions and for its “excessive reliance on non-binding ethical principles”, which may dilute enforceability and hinder uniform application.

On the American side, Gemini characterizes the Executive Orders as an “articulated and unified policy based on the principles of American global leadership, free-market innovation, and minimizing regulatory burdens,” which, in its view, “stands in stark contrast to the EU's human-centric and risk-based framework.” Among the core weaknesses identified, Gemini points to the overly “broad, non-specific policy goal of maintaining 'global AI dominance' without establishing clear metrics or safeguards,” which risks rendering the Orders aspirational rather than operational. Additionally, the Orders are critiqued for employing “highly subjective and politically charged definitions,” as well as reflecting an excessive “influence by private sector interests, potentially at the expense of other national objectives.” Of particular concern is what Gemini describes as an “aggressive and potentially legally vulnerable approach

to environmental regulations." According to this view, such measures are likely to provoke "intense legal and political opposition and may be challenged in court as an overreach of executive authority."

Mistral offers a structurally balanced assessment of the AI Act, viewing it as a regulation with "clear objectives", focusing on "improving the functioning of the internal market, promoting trustworthy AI, and ensuring high levels of protection for health, safety, and fundamental rights." Mistral commends the AI Act's "detailed provisions and annexes", especially in relation to general-purpose AI models and systemic risk assessment, which are seen as providing essential clarity and operational guidance. Nonetheless, Mistral identifies several limitations. Chief among these is the "complexity and length" of the regulation, which may present "challenges for implementation and compliance, particularly for smaller entities or startups"—raising concerns about proportionality and administrative burden. Additionally, Mistral flags "ambiguity in definitions", which may give rise to "inconsistencies in interpretation and application." In this respect, further definitional clarity and illustrative examples could enhance legal certainty. Finally, the absence of a clear strategy for "resource allocation" is seen as a significant shortcoming. Without specifying what support will be available to Member States and stakeholders, Mistral warns that the Act's effectiveness could be undermined at the point of implementation.

In contrast, Mistral views the Orders as programmatic blueprints, praised for articulating a "clear purpose and policy statement", which serves to establish a coherent directive orientation. In particular, the emphasis on "maintaining U.S. leadership in AI innovation" is viewed as central to the Orders' strategic ambition and geopolitical framing. Nevertheless, several limitations are noted. A recurring concern relates to "ambiguity in language," particularly in ideologically charged expressions

such as “free from ideological bias or engineered social agendas.” Mistral critiques such phrasing as “somewhat vague” and cautions that, absent clearer definitions or operational examples, these provisions may lead to uneven interpretation and problematic enforcement. Implementation is another area of uncertainty. Although the Orders set deadlines—such as the requirement to develop an action plan within 180 days—Mistral observes a lack of “detailed guidance on how this plan should be formulated or what specific benchmarks it should meet.” Similarly, the reference to federal resources, including public lands, is not accompanied by clarity on “the extent of these resources or the criteria for their allocation,” which may generate legal or interagency disputes. Mistral also flags the limited role assigned to non-governmental stakeholders, suggesting that the Orders would benefit from more explicit inclusion of “private sector entities, academic institutions, and international partners” to support a pluralistic and globally integrated regulatory model. This omission, in Mistral’s view, reflects a missed opportunity to align with collaborative governance norms already embraced in international AI frameworks. Furthermore, the absence of robust mechanisms for “accountability and oversight” is considered a structural weakness. Lastly, the Orders are critiqued for failing to engage with “potential ethical considerations or safeguards” in AI development and deployment—an omission that places the U.S. framework out of step with the emerging global consensus around rights-based and human-centric AI regulation.

From a more general point of view, when juxtaposing the two frameworks, all models tend to construct a polarity along the axes of speed versus thoroughness, and innovation primacy versus rights primacy. According to the qualitative interviews, most LLMs tend to read the EU Act as normatively “closed,” seeking *ex ante* certainty through comprehensive and rigid codification, while the US orders are more

normatively “open,” favoring policy agility and adaptive governance over exhaustive legal specification. This interpretative dichotomy extends to the treatment of contested normative terrain: for instance, the US prohibition on “DEI” considerations in AI procurement is consistently construed by the models as directly conflicting with the EU’s entrenched commitment to diversity and non-discrimination, thereby creating a transatlantic compliance tension of both legal and operational significance. Moreover, the LLMs’ views reveal an awareness of complementarities: the EU’s juridical density could be read as supplying rights-based safeguards absent in the US model, while the US’s industrial policy levers could compensate for the EU’s relative inattention to infrastructural acceleration and global competitiveness, reflecting the inherent design principles or training methodologies and amplifying biases attributable to their geographical origin[51] or the strategic imperatives of their developing entities[52].

A consistent point of analytical convergence across the LLMs pertains to the implementation challenges confronted by both regulatory texts, albeit of a qualitatively different nature. For the EU framework, the observed rigidity and inherent complexity of the AI Act, especially concerning its application to small and medium-sized enterprises (SMEs). This concern resonates with established critiques within European administrative law regarding the asymmetry between legal harmonization and institutional capacity[53]. In the US legal and administrative landscape, a primary impediment to the effective implementation of the Orders inheres in the fragmented

[51] Evidence in this regard can be found in Rohin Manvi, Samar Khanna, Marshall Burke, David Lobell and Stefano Ermon, 'Large Language Models are Geographically Biased' (2024) arXiv preprint < https://arxiv.org/abs/2402.02680> accessed August 10, 2025.

[52] Indicative in this respect is the human factor bias. See Rajesh Ranjan, Shailja Gupta and Surya Narayan Singh, 'A comprehensive survey of bias in llms: Current landscape and future directions' (2024) arXiv preprint <https://arxiv.org/abs/2409.16430> accessed August 10, 2025.

[53] This issue has been raised in a number of sectors and it is considered a critical weakness of the European Union, as shown in Franziska Bremus and Tatsiana Kliatskova, 'Legal harmonization, institutional quality, and countries' external positions: A sectoral analysis' [2020] Journal of International Money and Finance 1.

allocation of regulatory authority among diverse federal agencies. This jurisdictional dispersion is exacerbated by the discernible absence of a cohesive, centralized compliance framework. Consequently, this fragmentation and lack of centralization pose a material risk of compromising accountability and fostering an inconsistent application of prevailing legal norms, given that both regulations might be insufficiently "future-proof," for excessive specificity or insufficient definitions[54].

4. Deconstructing the AI Regulatory Paradox: A Discussion Towards Synergistic Governance Frameworks

The comparative analysis developed in this study reveals an apparent paradox at the heart of contemporary AI regulation: while both the EU and the US explicitly acknowledge the imperative of regulating artificial intelligence to preserve public trust, safeguard fundamental values, and mitigate systemic risks, they pursue these objectives through regulatory grammars that are not only divergent but structurally incommensurable. This divergence is not reducible to differences in institutional design; it reflects deeper ideological, constitutional, and epistemological commitments that shape each jurisdiction's understanding of the governance of emerging technologies. The AI Act embodies a model of legal determinacy rooted in a tradition of rights-based constitutionalism, precautionary risk management, and administrative legality. It adopts an *ex ante* rigid architecture, prescriptive in tone and systemic in logic, in which obligations are codified, risk categories are legally binding, and conformity assessments are embedded in a vertically integrated enforcement structure.

[54] This issue highlights a recurrent structural tension in contemporary regulatory theory: the inherent conflict between the desiderata of legal certainty and the profound challenge of maintaining technological responsiveness in the face of rapid and continuous innovation. The greatest difficulty lies in crafting regulations that remain pertinent and effective as technology relentlessly evolves. See Roger Brownsword and Han Somsen, 'Law, innovation and technology: before we fast forward – a forum for debate' [2009] Law, Innovation and Technology 1.

This approach reflects a juridical epistemology in which legal concepts are stabilized through codification and enforced through harmonized administrative mechanisms, producing high normative density and enforceability, yet also generating complexity and compliance asymmetries, especially for SMEs[55].

By contrast, the corpus of Executive Orders on AI governance advances an executive-led model that privileges institutional agility, sectoral innovation, and strategic policy direction. Its architecture is predominantly *ex post*, its tone aspirational, and its logic functional to an "anti-woke" agenda[56]. The Orders do not establish a comprehensive statutory framework; instead, they rely on a combination of agency discretion, inter-agency coordination, and policy signaling to guide AI development. Notably, they address objectives such as accelerating domestic AI innovation, ensuring "ideological neutrality" in federal AI procurement, promoting the export of US AI technologies, enhancing infrastructure deployment, and removing regulatory barriers. This managerial epistemology frames legitimacy not through codified obligations but through the responsiveness of institutions, the credibility of executive leadership, and the capacity to adapt governance to shifting technological priorities. While this enables flexibility and rapid policy recalibration, it also risks normative opacity, jurisdictional fragmentation, and reduced accountability[57].

The empirical findings substantiate these characterizations. The TF-IDF heatmap confirms the EU corpus's lexical concentration around terms such as *high-risk*, *conformity*, *provider*, *surveillance*, and *requirements*, all of which score markedly

[55] Critical arguments in this regard are pointed out in Thomas Joswig and Walter Kurz, 'Regulatory and Compliance Requirements for SMEs Operating AI Systems through Data Centers in the EU, with a Focus on Data Protection Challenges in Germany' [2025] Journal of Next-Generation Research 5.0 1.

[56] On this particular matter and for a more detailed discussion, see Dilys Schoorman, 'Waking up to the 'Anti-Woke'agenda' [2024] Journal of Educational Administration and History 404.

[57] For an analysis of the threats and opportunities associated with the new American legislator's approach, see Scott L. Greer, Holly Jarman, Rachel Kulikoff, Dimitra Panteli, Ewout van Ginneken and Matthias Wismar, 'The second Trump administration: A policy analysis of challenges and opportunities for European health policymakers' [2025] Health Policy 1.

higher than in the US corpus. This concentration indicates a governance model anchored in pre-defined legal categories and formalized compliance pathways. Conversely, the US corpus exhibits high salience for procedural and institutional terms such as *agency*, *order*, *appropriate*, *guidance*, and *shall*. The prominence of *shall* – with a TF-IDF weight of 0.71 compared to 0.36 in the EU corpus – signals the reliance on prescriptive directives within an otherwise non-codified, executive-driven framework. Furthermore, the semantic mapping and vector embeddings reveals that EU-specific terms cluster tightly around formal regulatory instruments, reflecting vertical integration of rulemaking, oversight, and enforcement. US-specific terms are more dispersed, reflecting a governance landscape reliant on distributed institutional action and policy guidance rather than cohesive statutory design. The comparative PCA 3D projections show minimal overlap in the semantic spaces occupied by the two corpora, underscoring the depth of their structural divergence. Finally, the qualitative interviews with LLM identify analogous deficiencies in both systems, including the under-specification of cross-border enforcement, the absence of integrated redress mechanisms, and ambiguity in the allocation of accountability in the AI supply chains.

Taken together, these empirical patterns corroborate the theoretical claim that neither the EU's uniform formalism nor the U.S.'s open-ended discretion suffices as a standalone model for AI governance. The over-determination of normative content in the EU framework risks inflexibility in the face of technological change; the under-determination of enforceable content in the U.S. framework risks governance gaps in areas demanding categorical protection. This evidence supports the necessity of a hybrid regulatory paradigm—one that synthesizes the EU's commitment to legal certainty, rights-based safeguards, and institutional accountability with the U.S.'s emphasis on innovation, procedural adaptability, and interagency coordination. Such a

model would need to balance rule-based clarity with dynamic responsiveness, embedding robust ethical and legal protections while preserving the agility required to govern rapidly evolving AI systems. Such a model would retain the EU's enforceable safeguards for high-risk applications while integrating the US's capacity for innovation-oriented flexibility, embedding both within a structure of procedural accountability, transparent experimentation, and transnational interoperability. In this way, the quantitative and semantic analyses do not merely describe divergence – they confirm the necessity of a principled synthesis that transforms structural asymmetry into a resource for adaptive and proportionate governance.

5. Conclusion

The comparative investigation conducted throughout this paper has revealed a significant divergence between the regulatory architectures adopted by the EU and the US in their respective approaches to artificial intelligence governance. This divergence is not limited to differences in institutional design or procedural implementation; rather, it reflects fundamentally different conceptions of the role of law in structuring digital markets, managing systemic risk, and ensuring accountability in the development and deployment of AI technologies. While the EU approach is characterized by a detailed, rule-based framework aimed at harmonization and ex ante compliance, the US model relies more heavily on interagency coordination, soft law instruments, and market-driven innovation incentives. These contrasting models raise important questions regarding regulatory effectiveness, legal certainty, and the allocation of compliance burdens across public and private actors operating in transnational AI markets.

The EU's Artificial Intelligence Act establishes a normative model of regulation grounded in rights-based safeguards and categorical legal obligations. Its legal rationality is ex ante and preventive, structured around the codification of risk thresholds, the formal attribution of responsibilities, and the articulation of procedural guarantees for both compliance and redress. It is conceived as a proactive instrument of systemic control, designed to anticipate, classify, and constrain technological development before harm occurs. In this model, legality is expressed through structured anticipation and regulatory certainty, with the aim of harmonizing standards across Member States and mitigating risks at the design stage. By contrast, the US framework reflects a post hoc, executive-led approach centered on interagency coordination, policy signaling, and flexible administrative discretion. Rather than imposing binding obligations *ex ante*, it relies on non-binding instruments that are inherently adaptive and responsive. In this model, regulatory interventions emerge through institutional negotiation and political calibration, rather than through codified legal norms. The US approach thus leverages a form of "steering-based governance"[58], in which law operates primarily as an infrastructural mechanism to facilitate innovation, workforce development, and national competitiveness. Here, the function of legal norms is not to constrain public and private power through enforceable limits, but to enable policy implementation through coordination, incentives, and strategic alignment across federal agencies and industry actors.

The structural divergence between the two legal frameworks across the Atlantic generates what this paper has called as an "AI regulatory paradox": both systems aim to secure safety, trust, and accountability in AI, yet they do so through irreconcilably

[58] Cf. Kaisa Lähteenmäki-Smith, Samuli Manu, Pirkko Vartiainen, Petri Uusikylä, Harri Jalonen, Sami Kotiranta, Urho Lintinen, Mikko Annala, Iacopo Gronchi, Juha Leppänen and Silva Mertsola, *Government Steering Beyond 2020: From Regulatory and Resource Management to Systems Navigation* (Finnish Publications of the Government's Analysis, Assessment and Research Activities | 2021)

different institutional logics and normative assumptions. While the European model risks ossification – i.e., an excess of formalism that may hinder technological experimentation or lag behind rapid innovation cycles[59] – the American model risks evanescence, in that its normative aspirations might lack enforceable substance, undermining its capacity to deliver meaningful accountability or rights protection[60].

The computational methods employed – word clouds, TF-IDF vectorization, semantic embedding, and interpretive triangulation through state-of-the-art large language models – have rendered this transatlantic divergence empirically observable. The spatial and semantic separation between EU and U.S. legal vocabularies, the identification of conceptual outliers and partial overlaps, and the latent regulatory gaps revealed through machine-assisted analysis collectively confirm the absence of a shared juridical language – either in form or substance – through which AI can currently be governed in a coherent and enforceable manner across both jurisdictions. Nonetheless, the presence of a thin layer of shared regulatory vocabulary – evident in recurring terms such as *risk*, *accountability*, *transparency*, *trust*, and *governance* – indicates that, despite formal and structural divergences, both legal orders converge on a common set of foundational regulatory concerns. This suggests the possibility of developing forms of coordinated legal pluralism, whereby heterogeneous legal systems may functionally align on key regulatory objectives without necessitating the collapse into a single harmonized framework.

Against this backdrop, this paper advocates for the construction of a risk-based adaptive regulatory architecture, capable of modulating legal obligations, institutional density, and procedural rigor in proportion to the systemic criticality and epistemic

[59] On legal ossification see, among others, Ran XI, 'On Emerging Technologies: The Old Regime and the Proactivity' [2025] Cardozo International and Comparative Law Review 75.
[60] See, in this regard, the critical notes emerging from Brian Judge, Mark Nitzberg and Stuart Russell, 'When Code Isn't Law: Rethinking Regulation for Artificial Intelligence' [2025] Policy and Society 85

opacity of the AI systems subject to regulation. Rather than choosing between rigidity and flexibility, this model seeks to organize their relationship: it draws on the granularity and enforceability characteristic of the European legal framework for high-risk AI, while incorporating the institutional agility and reflexivity emblematic of the American system in innovation-driven contexts where anticipatory regulation risks stifling technological progress or proving ineffective under conditions of epistemic uncertainty. In this framework, risk should not be regarded as a neutral descriptor but rather as a legal operator and vector for the allocation of responsibility, regulatory obligations, and oversight. Implementing such a model requires not only the calibration of normative intensity according to differentiated risk profiles but also the institutionalization of regulatory learning processes, transparency by design, and mechanisms ensuring reversibility and contestability.

General-purpose AI systems, foundation models, and cross-border data flows generate jurisdictional entanglements that challenge traditional conceptions of regulatory sovereignty. Consequently, adaptive regulation must also prioritize interoperability: enabling reciprocal recognition of audit mechanisms, establishing shared principles of algorithmic accountability, and enforcing minimum thresholds of institutional oversight. What is required is not the creation of a supranational legal order, but rather the development of a coordinated architecture of differentiated yet communicable norms, through which legitimacy, enforcement, and interpretability can be effectively distributed across multiple legal systems without succumbing to fragmentation or inconsistency.

This paper's contribution is therefore twofold: it provides an empirical demonstration of the semantic and institutional divergence between the two leading AI regulatory systems, and it offers a juridically grounded proposal for transforming that

divergence into a generative resource. It can be argued that legal scholars, regulators, and technologists alike must move beyond binary frameworks of harmonization versus fragmentation and begin designing legal infrastructures that are responsive to varying degrees of risk, depths of uncertainty, and differentiated expectations of legality. In short, legality in the age of AI must be reconceptualized as structured adaptability, governed by principles of proportionality, epistemic humility, and institutional interdependence. Only by integrating the respective strengths of both paradigms—the European emphasis on safeguarding fundamental collective rights and the American focus on preserving space for innovation—can the law meaningfully respond to the challenges posed by learning systems that increasingly govern, decide, and act in domains once reserved exclusively to human and public judgment.