

Book Review

A Review of Matthijs M. Maas' Book, *Architectures of Global AI Governance*

Matthijs M. Maas, *Architectures of Global AI Governance: From Technological Change to Human Choice*. Oxford: Oxford University Press, 2025. ISBN 978-0-19-887783-7

Beatriz Rayón-Viña^{1,*}

¹ University of Oviedo 1; rayonbeatriz@uniovi.es

* Correspondence

Abstract: This review examines Matthijs M. Maas' work *Architectures of Global AI Governance: From Technological Change to Human Choice*. The context in which this work is presented is the qualitative leap in the development of AI-based systems since 2016, and especially since 2022, when the pursuit of general-purpose systems intensified amid talk of an arms race. The fundamental question running through the book is how to govern a highly variable technology (where outcomes and risks are extremely context-dependent) that evolves much faster than law and diplomacy. To answer this, Maas constructs a framework of three complementary lenses: the first reveals that regulation fails when it focuses on the artifact rather than on the sociotechnical changes it enables; the second analyzes how that same AI disrupts the legal tools with which we attempt to govern it; and the third describes the actual institutional landscape: not a single, large governing institution, but a polycentric and fragmented ecosystem to which AI governance must adapt and which, in turn, it will reshape. It is this articulation between levels (sociotechnical, legal, and institutional) that serves as the framework from which this review will assess both its strengths and its limitations.

Keywords: artificial intelligence; governance; sociotechnical change; governance disruption; polycentric institutions; emerging technologies; global governance

Citation: Rayón-Viña, Beatriz. 2026.

A review of Matthijs M. Maas' book

"Architectures of Global AI

Governance". *Journal of Ethics and*

Emerging Technologies 36: 2.

<https://doi.org/10.55613/jeet.v36i2.240>

0

Received: 17/05/2026

Accepted: 18/05/2026

Published: 01/07/2026

Publisher's Note: IEET stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The work is presented in the context of the developmental leap that has occurred in the development of AI-based systems since 2016 (when the era of large-scale AI began) and especially since 2022, when the quest for general-purpose AI systems began, framed within the narrative of an arms race. In this context, the first key point to grasp from the book is the complexity involved in monitoring a Large Language Model (LLM). This challenge is what, according to Maas, constitutes a blurring of the traditional boundaries of law. The fundamental question is how we can regulate and legislate a technology that evolves much faster than law and diplomacy.

If AI-based systems are presented as a new geopolitical power, the battle, Maas tells us, no longer involves only the U.S.-China fronts; rather, "challenger" states such as the United Arab Emirates, Saudi Arabia, France, South Korea, or Singapore also come into focus: powers that are investing heavily in their own models because they seek sovereignty. The fundamental characteristic of AI is its high variability. In other words, it is not a monolithic technology with uniform performance, but an extremely heterogeneous field where remarkable successes coexist with significant failures:

outcomes and risks are extremely diverse and context-dependent, which requires flexible risk management rather than a simplistic view of the technology.

In order to address this issue, Maas presents an integrated analytical framework comprising three conceptual lenses that help us understand and design the governance of AI-based systems in today's environment, which I will outline and analyze in this review.

2. Regulating problems, not artifacts: the socio-technical turn

From the perspective of STS, it is essential to consider artifacts as an analytical lens through which one can reconstruct networks of actors, institutions, norms, knowledge, and conflicts that would otherwise remain invisible. In this triple vision of AI governance that Maas offers, the author proposes going beyond the regulation of the artifact itself and highlights the need to shift from artifact-centered regulation to regulation oriented toward sociotechnical changes. Techno-centric approaches are obsessed with the "novelty" of the algorithm or model and tend to overemphasize what is exceptional about the technology, whilst underestimating the continuity with already known problems of market power, inequality, or regulatory capture. Similarly, this problem arises in legislative frameworks, which are fragmented by application and address visible symptoms rather than the structural patterns that AI reproduces across sectors and jurisdictions.

The book suggests viewing AI as a "sociotechnical stack" composed of four factors: (1) computational techniques (the set of algorithms and methods that improve the accuracy and speed of machine decision-making), (2) capabilities (the functions produced by these techniques that can complement or replace human performance), (3) applications (the use of these capabilities in specific domains and industries), and (4) sociotechnical impacts (the final level where these applications generate changes in human behavior, social relations, and power structures) (Maas, 2025, p.302). In this approach, the subject of regulation is not the line of code, but rather the transformations in social organization that this technology enables or accelerates. It is here that Maas finds the source of his position's legitimacy: he establishes that the need to intervene arises when certain critical configurations occur, and the author presents us with a taxonomy of six "logics of the problem": ethical challenges, security threats, physical safety risks, structural changes, common benefits, and disruption in governance (Maas, 2025, p.347).

The most interesting contribution of this framework is the question of what kind of problem lies before us. Rather than compiling generic lists of principles, the regulator can ask what logic prevails in a specific case and, based on that, which legal or institutional instruments are most appropriate. However, the sophistication of the framework also highlights one of its potential limitations: it requires analytical capabilities and cross-sectoral coordination that many public administrations lack, especially in contexts with a shortage of specialized staff and strong political pressures. To this end, in Chapter 7 the author addresses the issue of epistemic communities (networks of individuals or groups possessing an authoritative claim to knowledge relevant to policy within their domain of expertise) and governance experts, who can serve as a bridge to facilitate this effective communication (Maas, 2025, p.523). However, a potential shortcoming of this approach is that the author does not explore in detail how these capabilities are built or who funds and oversees this "translation" work between sociotechnical theory and regulatory decision-making. That said, this shift toward sociotechnical change as a regulatory object has an immediate consequence: if AI alters social structures, it also alters the legal tools with which we attempt to regulate it, which brings us to the second lens.

3. Disruption of governance: when AI disrupts the legal system itself

The second lens presents AI as a force that reshapes the very tools of governance. The “disruption of governance” discussed in Chapter 5 describes situations in which technology not only creates new borderline cases but also undermines the substance of the law, its processes of creation and application, and the political foundations that give it its legitimacy (Maas, 2025, p.377). AI is presented as a catalyst for “constitutional” transformations in the broadest sense: it redefines what counts as evidence, how intent or authorship is attributed, and what forms of delegation to automated systems are considered admissible.

Three main dynamics of disruption can be identified: in the development of governance, where AI creates legal loopholes or gray areas that force the law to evolve; the displacement of governance, when AI automates or replaces established legal practices; and, finally, the destruction of governance, where AI contributes to the erosion or decline of existing regulatory regimes—perhaps the most important aspect of this second lens. According to Maas, this is due to the law’s inability to adapt dynamically (blocked by political interests or conceptual friction, as is the case with attempts to regulate cyberweapons lacking clear physical signatures). On the other hand, this can also intensify when AI is used as an “anti-law weapon”: tools designed to manipulate public opinion, undermine trust in electoral processes, or systematically exploit the weaknesses of multilateral institutions.

Perhaps the most problematic aspect of this approach is the inherent difficulty of designing technologically neutral laws, as it requires striking a balance between regulatory generality (to avoid obsolescence, but which can become overly broad and thus inefficient) and procedural precision (to ensure that the law can be interpreted and applied consistently as technology evolves). Lastly, Maas’ analysis is less detailed when it comes to explaining what incentives actors who benefit from opacity or the speed of technological change would have to voluntarily accept these self-restraint mechanisms. Here, the disruption of governance is described with great diagnostic clarity, but the dimension of political and power conflict (who gains and who loses with each reform) is treated relatively tangentially. Yet the disruption of the law does not occur in a void, but within an institutional ecosystem already saturated with overlapping regimes; understanding that complexity is the step that completes the framework

3. The complexity of systems: governing in a chaotic institutional ecosystem

The third lens focuses on the global institutional environment in which AI governance must operate. The book adopts the concept of regime complexity to describe an architecture in which, within a single thematic area, multiple institutions, treaties, and forums operate with overlapping and sometimes contradictory agendas. In contrast to the centralist solution of creating a large, overarching governing institution, the author argues that the most realistic scenario is a polycentric, dense, and fragmented ecosystem into which the new governance of AI must fit and which, in turn, it will help to reshape. This is justified by the current existence and plurality of regulatory frameworks and specialized institutions that already constitute a “global institutional ecology.” In this context, AI will not arrive in a blank slate, but rather in a fabric already saturated with regimes governing human rights, trade, security, telecommunications, civil aviation, the environment, and cybersecurity, among others. Interactions between these regimes can produce gaps, regulatory conflicts, loose cooperation, or mutually reinforcing synergies, and it is precisely this tangle that makes up the real playing field of AI governance.

In response to this proposal, there are two clear positions: on the one hand, critics argue that institutional fragmentation leads to regulatory conflicts, encourages forum shopping by the most powerful actors, and widens inequalities between states. On the other hand,

its defenders argue that polycentric systems are more flexible and resilient because they allow for the testing of diverse responses, reduce the risk of capture by a single institution, and open space for regulatory innovations from the periphery. This is the author's position, who justifies the need for this plurality on the grounds that such systems are inherently more flexible and capable of adapting to the rapid pace of technological change in AI.

Finally, the book summarizes this approach in a five-step framework for analyzing and, to the extent possible, guiding the complex landscape of AI regimes: examine the origins and political viability of new regimes, map the existing institutional landscape, track trends toward greater fragmentation or integration, assess the consequences in terms of effectiveness and legitimacy, and design strategies to strengthen the overall effectiveness, resilience, and coherence of the system (Maas, 2025, p.557-558). The proposal invites us to stop searching for "the best institution" in the abstract and instead to consider how different components (treaties, technical bodies, private standards, ad hoc coalitions) can be coordinated to manage a technology that, by design, transcends any single silo.

4. Conclusion

Finally, it is worth noting a formal observation regarding the manuscript. The work is quite lengthy, which does not always work in its favor: the first chapters, devoted to contextualizing the development of AI and justifying the urgency of its governance, have a more expository than analytical tone, which delays the reader's entry into the book's core argument. It is not until Chapter 4 that the three-lens framework begins to unfold in full force, and where the reading becomes genuinely stimulating. This imbalance between the introductory section and the main body of the analysis creates a certain sense of discontinuity, as if the volume had combined two writing projects with distinct logics. Added to this is a proliferation of subheadings that, far from facilitating navigation through the text, at times unnecessarily fragment the argumentative thread. None of these observations invalidates the solidity of the proposed framework, but it does suggest that a more restrained edition would have better concentrated the impact of its most original contributions.

In any case, taken together, the three lenses offer more than just a mosaic of concepts; they constitute an attempt to reframe the conversation on AI governance. The way it articulates levels (sociotechnical, legal, and institutional) is one of the book's greatest strengths, as it avoids both technodeterminism and naive legalism. The book provides a clear conceptual map from which to formulate the difficult questions by shifting the focus toward sociotechnical change, the disruption of governance, and the complexity of regimes. The author suggests that seriously discussing AI regulation necessarily involves talking about priorities, the limits of the law, and institutional architecture—not just lists of well-intentioned ethical principles. However, work remains to be done in translating this theoretical framework into empirical research agendas and viable negotiation strategies in a deeply unequal world. The window of opportunity to regulate AI is not infinite, and precisely for that reason, Maas argues, it matters to have frameworks that combine conceptual ambition with a more direct focus on power struggles and the actual capabilities of the actors tasked with governing it.

Acknowledgements: The original manuscript was drafted in Spanish by the author. DeepL was used for translation into English. Perplexity AI was consulted for stylistic and structural refinement of the translated text. All analysis, critical assessment, and scholarly argumentation are the sole responsibility of the author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

Maas, M. M. (2025). *Architectures of global AI governance: From technological change to human choice*. Oxford University Press.
<https://doi.org/10.1093/9780191988455.001.0001>