Book Review of *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence*

Beyond Hype: What the Handbook Offers for Responsible AI Governance

Oihika Arpit
Graduate Researcher, upGrad
Kshitij Taluja
Undergraduate Researcher, Bachelor of Computer Science, Golden Gate University
Dr. Smrite Goudhaman
Global Professor of Practice, Golden Gate University

Editor: Nathalie A. Smuha (Ed.)
Title: *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence*
Publisher: Cambridge University Press
Year: 2025
ISBN: 978-1-009-36781-3
Pages: 460

Abstract

The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence, edited by Nathalie A. Smuha, is most useful when read not as another catalogue of AI risks, but as an argument about how societies should keep human judgment, rights, and institutional responsibility in view as algorithmic systems become ordinary parts of decision-making. Its three-part architecture moves from technical and philosophical foundations to law and policy, and then to applications in education, media, healthcare, finance, employment, policing, public administration, and armed conflict. The volume's strongest contribution is the way it connects questions that are too often separated: what a system can technically do, what values it embodies, what legal rules constrain it, and who remains answerable when things go wrong. Its discussions of design for values, fairness, sustainability, the GDPR, and the EU AI Act are particularly effective in exposing the distance between principles and implementation. The book is not without limits. Its European legal focus, specialist vocabulary, and edited-chapter format mean that readers looking for a globally comparative or step-by-step implementation manual will need additional sources. Even so, the *Handbook* offers a mature starting point for engineers, policymakers, educators, researchers, and organizational leaders who want to approach AI as a social and institutional technology rather than merely a computational one.

Keywords: Artificial intelligence; AI governance; AI ethics; law and technology; human rights; EU AI Act; algorithmic systems; responsible innovation

Artificial intelligence is often discussed as though the main question were how quickly the technology can improve. A second conversation then begins, usually after the excitement: how quickly can regulation catch up? *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* takes a more useful route. It asks what happens when systems built to predict, classify, recommend, generate, or decide become woven into institutions that already distribute rights, opportunities, resources, and responsibilities. The book's answer is deliberately plural. Ethics, law, policy, institutional practice, and public accountability have to meet at the point where an AI system actually affects a person.

That is also why the *Handbook* feels different from a conventional AI textbook. Edited by Nathalie A. Smuha, the volume brings together computer science, philosophy, ethics, law, social science, and sector-specific expertise. Part I, “AI, Ethics and Philosophy,” establishes the technical and conceptual ground; Part II, “AI, Law and Policy,” asks what existing and emerging legal instruments can do; and Part III, “AI across Sectors,” tests these ideas against real institutional settings. The sequence gives the book an underlying argument even though the chapters are written by different authors. The question this review returns to is whether that interdisciplinary argument ever becomes guidance an institution can actually act on.

Part I is particularly strong because it refuses to let ethics float above the engineering. The chapter on ethics argues for a “Design for Values” approach, in which ethical considerations are incorporated into the design process rather than attached as a late-stage checklist (Buijsman et al., 2025). It suggests that choices about data, objectives, system architecture, interfaces, deployment, and monitoring are never entirely neutral. The practical implication is not that engineers should become philosophers, but that responsible engineering requires the ability to recognize when a technical decision is also a decision about people and institutions. What the chapter says less about is what happens when those values collide, or when the stakeholders around a system disagree about which value should win.

The discussion of fairness develops the same point with greater conceptual precision. Naudts and Vedder (2025) distinguish procedural and substantive approaches to fairness and place them alongside justice, equality, and social relationships. Their warning against treating fairness as a problem that can simply be solved through model design is important. A model may satisfy a selected statistical criterion and still operate inside an unfair institution. This is one of the *Handbook*'s clearest objections to techno-solutionism: improving the model does not necessarily repair the process in which the model is used. The chapter therefore

encourages a broader question: fair according to whom, within which social arrangement, and with what consequences? For decision-makers, the implication is uncomfortable: fairness has to be audited at the level of the process, not certified at the level of the model.

The chapter on moral responsibility adds another layer. Lauwaert and Oimann (2025) examine whether autonomous technologies create a genuine "responsibility gap" and resist the easy assumption that autonomy automatically removes human responsibility. The *Handbook* instead keeps attention on the human and institutional arrangements surrounding an AI system. The harder question is left open. When developers, vendors, deployers, managers, and end users all shape how a system behaves, the Handbook says less about how responsibility should be divided among them. The chapter on sustainability raises a related concern. Hasselbalch and Van Wynsberghe (2025) show that AI can be used in support of sustainable development while also carrying environmental and social costs of its own.

Part II shifts from normative questions to the legal structures through which AI is governed. The GDPR chapter shows how principles such as lawfulness, transparency, purpose limitation, data minimization, and accountability can come under pressure when systems depend on large-scale data processing (Dewitte, 2025). The chapters on tort liability, competition, consumer protection, and intellectual property widen the picture. The cumulative effect is persuasive: AI governance is a legal ecosystem, not a single statute.

The discussion of the European Union AI Act is arguably the book's most important test of its own governing philosophy. Smuha and Yeung (2025) acknowledge the Act's ambition to establish a comprehensive framework while examining whether its architecture can deliver the protection its objectives promise. Their approach is especially useful because it does not equate the existence of regulation with the existence of effective governance. A rule may express the right values and still fail if oversight is weak, responsibilities are unclear, or affected people have limited avenues for challenge. This makes the analysis more credible than the familiar "Europe has regulated AI" narrative. The chapter is stronger on what the Act promises than on how it will be enforced, which is where the gap between principles and practice usually opens.

The *Handbook* becomes more concrete in Part III. Here the reader sees how the same technology can carry different meanings and risks depending on where it is deployed. In education, for example, Molenaar et al. (2025) distinguish replacement from augmentation and emphasize forms of AI that support rather than simply displace learners and teachers.

They also present the idea of embedded ethics and collaborative development involving educational professionals, researchers, and technology providers.

The sectoral chapters also prevent the reader from treating “AI risk” as a single, portable category. Predictive policing makes a particularly sharp intervention because accuracy is not the only relevant test; the legitimacy of the deployment and its effects on citizens matter as well. Public administration raises an even broader concern: when governments delegate parts of public decision-making to algorithmic systems or private technology providers, efficiency can come into tension with digital sovereignty, democratic accountability, and the rule of law. Across these settings the risks change but the governance problem does not: someone must remain answerable, oversight must be real, and the people affected need a way to object.

The final sectoral chapter, on armed conflict, confronts the hardest version of these questions. It considers whether AI and other digital technologies can support compliance with international humanitarian law while asking whether existing legal frameworks are adequate for the technologies now entering military contexts. AI can become part of decisions involving life, liberty, public power, and armed violence. The scale of the stakes makes the book's insistence on accountability harder to dismiss as a theoretical preference.

What ultimately holds these sections together is the concluding argument that law is necessary but not sufficient. Smuha's concluding remarks emphasize the continuing relevance of existing legal rules, while also acknowledging gaps, enforcement problems, and tensions within the legal patchwork (Smuha, 2025). The book points toward AI literacy, education, responsibility-taking, public deliberation, awareness, and critical engagement as complementary forms of governance. This is perhaps the *Handbook*'s most durable contribution. Responsible AI is presented less as a compliance box to tick than as a capacity that institutions have to build and maintain.

The book does, however, have boundaries that readers should understand. Its primary legal and normative frame is European. That focus is entirely defensible given the expertise of the contributors and the importance of European AI governance, but it limits the book's ability to explain how the same principles operate in legal and institutional environments outside Europe. Regulatory capacity, enforcement machinery, data-governance structures, and the way rights are framed all differ across jurisdictions, so the same principle can produce very different practice. Readers working in South Asia, Africa, Latin America, or other governance contexts will need comparative material alongside it. Finally, the edited format

produces variation in method and emphasis. The reward is a wide range of perspectives; the cost is that the reader does not receive one operational framework that can be lifted directly into an organization's AI governance process.

Those limitations actually help clarify the book's proper use. This is not a coding manual, a product-selection guide, or a corporate compliance checklist. It is better understood as a governance and critical-thinking handbook. For engineers, it provides language for seeing how technical choices can become questions of rights, fairness, sustainability, and legitimacy. For organizational leaders, perhaps the most practical lesson is that responsible adoption also requires decisions about where AI should not be used, who retains responsibility, what evidence is sufficient for deployment, and how people can contest consequential decisions. On balance, the missing operational framework is a boundary rather than a flaw. Building one is not the job of a multidisciplinary handbook, though readers should know they will have to look elsewhere for it.

A senior handbook should leave a reader with better questions, not merely more information. This one does. After reading it, it becomes harder to ask only whether an AI system works. One is pushed to ask what “works” is supposed to mean, whether oversight is meaningful, and what happens when the system behaves differently from what its designers expected.

Overall, *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* is a substantial and timely contribution to the responsible-AI literature. Its greatest strength is its refusal to separate technological capability from social consequence. The book does not offer a universal recipe for responsible AI, and it would be misleading to pretend that it does. What it offers instead is more durable: an interdisciplinary way of examining AI through the combined lenses of values, law, institutions, and human agency. The *Handbook*'s central message is ultimately a modest but demanding one: society should not allow the pace of technological capability to determine the limits of human accountability.

References

Buijsman, S., Klenk, M., & van den Hoven, J. (2025). Ethics of AI: Toward a “Design for Values” approach. In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 59–78). Cambridge University Press. https://doi.org/10.1017/9781009367783.005

Dewitte, P. (2025). AI meets the GDPR: Navigating the impact of data protection on AI systems. In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and*

Policy of Artificial Intelligence (pp. 133–157). Cambridge University Press. https://doi.org/10.1017/9781009367783.010

Hasselbalch, G., & Van Wynsberghe, A. (2025). Artificial intelligence, power and sustainability. In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 117–130). Cambridge University Press. https://doi.org/10.1017/9781009367783.008

Lauwaert, L., & Oimann, A.-K. (2025). Moral responsibility and autonomous technologies: Does AI face a responsibility gap? In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 101–116). Cambridge University Press. https://doi.org/10.1017/9781009367783.007

Molenaar, I., Baten, D., Bárd, I., & Stevens, M. (2025). Artificial intelligence and education: Different perceptions and ethical directions. In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 261–282). Cambridge University Press. https://doi.org/10.1017/9781009367783.017

Naudts, L., & Vedder, A. (2025). Fairness and artificial intelligence. In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 79–100). Cambridge University Press. https://doi.org/10.1017/9781009367783.006

Smuha, N. A. (Ed.). (2025). *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence*. Cambridge University Press. https://doi.org/10.1017/9781009367783

Smuha, N. A., & Yeung, K. (2025). The European Union's AI Act: Beyond motherhood and apple pie? In N. A. Smuha (Ed.), *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence* (pp. 228–258). Cambridge University Press. https://doi.org/10.1017/9781009367783.015