



INTERNATIONAL JOURNAL OF HUMAN RIGHTS LAW REVIEW

An International Open Access Double Blind Peer Reviewed, Referred Journal

Volume 5 | Issue 5 | 2026

Art. 11

Beyond the Hype: Artificial Intelligence, Institutional Failure and the Limits of Prediction

Saksham Bharadwaj

Law Student, BBA.LL.B., 4th Year,

Himachal Pradesh National Law University, Shimla

Recommended Citation

Saksham Bhardwaj, *Beyond the Hype: Artificial Intelligence, Institutional Failure and the Limits of Prediction*, 5 IJHRLR 117-126 (2026).

Available at www.ijhrlr.in/current-issues/.

This Article is brought to you for free and open access by the International Journal of Human Rights Law Review by an authorized Lex Assisto & Co. administrator.

For more information,
please contact submission.ijhrlr@gmail.com

Beyond the Hype: Artificial Intelligence, Institutional Failure and the Limits of Prediction

Arvind Narayanan and Sayash Kapoor, AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference, Princeton University Press, Princeton, 2024, 1st edn., 360 pp., ISBN 9780691249131.

INTRODUCTION: WHY AI NEEDS A LEGAL REALITY CHECK

Artificial intelligence is increasingly presented not merely as a technological development but as a force capable of transforming institutions, labour markets, public administration and democratic life. The legal significance of that transformation lies less in the abstract question of whether machines will become “intelligent” than in the concrete question of who is permitted to deploy automated systems, on what evidence, against whom, and with what safeguards. Arvind Narayanan and Sayash Kapoor’s *AI Snake Oil: What Artificial Intelligence Can Do, What It Can’t, and How to Tell the Difference* enters this debate at an important moment. Published by Princeton University Press in 2024, the book distinguishes between different forms of AI and challenges the tendency to treat the label “AI” as a guarantee of technical competence. Its central warning is that institutional actors can be persuaded to purchase or deploy systems that appear scientific while being unable to perform the tasks for which they are marketed.¹

The book deserves attention from legal scholars because this is ultimately a problem of governance. A prediction system used in recruitment, insurance, credit, policing or criminal justice does not remain a technical instrument once its output affects a person’s liberty, livelihood, equality or access to essential services. The authors’ scepticism therefore intersects with familiar legal concerns: non-discrimination, due process, privacy, accountability, transparency and the allocation of responsibility between developers and users. Their analysis is particularly useful because it moves the debate away from an undifferentiated question—whether AI is “good” or “bad”—towards the more legally productive questions of what a system can actually do and whether the institution using

¹ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can’t, and How to Tell the Difference* 1–35 (Princeton University Press, Princeton, 2024).

it is capable of governing its consequences.

This review argues that *AI Snake Oil* is strongest when it exposes the institutional demand for

unreliable prediction and insists on distinguishing technological capability from marketing claims. Its principal limitation, from a legal perspective, is that it does not develop a sufficiently systematic theory of rights, remedies or regulatory institutions. The book is consequently best understood not as a complete framework for AI regulation, but as an important diagnostic work that gives legal scholars a stronger factual and conceptual basis for designing regulation.

CENTRAL THESIS AND INTELLECTUAL FRAMEWORK

The book's intellectual framework rests on a basic but consequential proposition: "artificial intelligence" is an umbrella term covering technologies with substantially different capabilities and limitations. The authors distinguish predictive AI from generative AI and examine why systems that perform pattern recognition or generate plausible content should not automatically be treated as reliable instruments for forecasting human behaviour.² This distinction is foundational to the book's argument because legal rules that regulate "AI" as a single category risk either over-regulating relatively benign applications or under-regulating high-impact systems.

The authors' principal target is what they call "AI snake oil": systems that do not and cannot work as advertised.³ Their criticism is not equivalent to the claim that AI is useless. On the contrary, the book recognises genuine applications in which machine learning can assist with difficult tasks, particularly where large quantities of data can reveal patterns that are difficult for humans to identify. The problem arises when institutions ask predictive systems to answer questions for which the available data cannot provide a stable or meaningful basis for prediction.

The argument develops through several thematic stages. The opening chapters establish the distinction between types of AI and examine how predictive systems can fail. The discussion then considers the structural limits of predicting the future, the development of generative AI, the debate over existential risk, the possibility of using

² Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference* 36-98 (Princeton University Press, Princeton, 2024).

³ *Id.* 1-35.

AI to solve problems associated with social media, and the persistence of AI myths. The final chapter turns towards what society should do with this technology.⁴ This progression is significant: the authors begin with technical distinctions, move to institutional and social consequences, and ultimately arrive at a governance question.

The methodological character of the book is interdisciplinary rather than doctrinal. It combines explanations of machine learning with case studies, institutional examples, discussions of scientific evidence and observations about incentives within companies, research communities and the media. This approach makes the book accessible to lawyers who may not have technical training. At the same time, the non-doctrinal methodology means that the book is not designed to answer legal questions in the manner of a conventional law treatise. Its contribution is therefore upstream of legal doctrine: it helps determine what factual assumptions law should and should not make about AI.

THEMATIC STRUCTURE: FROM PREDICTIVE FAILURE TO INSTITUTIONAL RESPONSIBILITY

Rather than proceeding as a conventional technical textbook, *AI Snake Oil* builds a cumulative argument. The chapter on predictive AI examines failures in contexts where organisations seek to infer future outcomes from historical data. The following chapter explains why prediction is not simply a matter of feeding more data into a sufficiently powerful model. The problem may be structural: human behaviour is shaped by changing social conditions, institutional decisions and feedback effects, while the historical data on which a model is trained may encode the very inequalities the institution seeks to overcome.⁵

This insight has direct relevance to legal decision -making. If an algorithm predicts that an applicant is likely to default, reoffend or perform poorly at work, mathematical presentation does not establish neutrality. The data, variables and objective may reproduce existing inequalities. The authors' analysis therefore complements scholarship showing that algorithms can create discrimination risks while also offering opportunities for greater transparency when their design and operation are documented.⁶

⁴ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*, Contents vii–x (Princeton University Press, Princeton, 2024).

⁵ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference* 36–98 (Princeton University Press, Princeton, 2024).

⁶ Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan and Cass R. Sunstein,

The discussion of generative AI is comparatively cautious. The authors distinguish large language models from predictive systems used for high-stakes decisions and recognise that generative AI can be useful while remaining unreliable. This matters to lawyers because fluent outputs can conceal factual or legal errors and encourage users to attribute more authority to them than is justified.

The later chapters widen the analysis beyond technological capability by asking why AI myths persist. The explanation points towards incentives: vendors may exaggerate capabilities, institutions may seek technological shortcuts, and research and media environments may reward novelty.⁷ This prevents technological failure from being treated as merely an engineering problem; “AI snake oil” persists because it satisfies institutional demand.

CRITICAL APPRAISAL: STRENGTHS OF THE ARGUMENT

The first major strength of *AI Snake Oil* is conceptual clarity. Public debate frequently collapses predictive models, recommendation systems, facial recognition, large language models and hypothetical artificial general intelligence into a single category. By separating these technologies, the authors provide a more rational basis for legal analysis. A regulator cannot meaningfully assess risk without identifying the function a system performs, the population affected, the degree of autonomy involved and the consequences of error.

The second strength is the authors’ emphasis on evidence rather than technological prestige. The book asks readers to interrogate the evidentiary foundation of claims made about AI. This is particularly valuable in law, where the appearance of scientific objectivity can influence courts, administrative bodies and policymakers. A statistical output should not be treated as inherently more reliable merely because it is produced by a computer. The legal question remains whether the evidence is sufficiently valid, relevant and contestable for the decision being made.

Third, the book identifies a subtle but important relationship between AI and institutional failure. Institutions may turn to automated systems because existing human processes are expensive, inconsistent or overwhelmed. The promise of prediction therefore

“Discrimination in the Age of Algorithms”, 10 *Journal of Legal Analysis* 113, 113 – 174 (2018).

⁷ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can’t, and How to Tell the Difference* 227–257 (Princeton University Press, Princeton, 2024).

becomes attractive precisely where governance is already weak.⁸ This observation helps explain why replacing human discretion with an algorithm does not necessarily remove arbitrariness. It may merely relocate discretion into the selection of training data, the design of the model, the definition of success and the institutional decision to rely upon its output.

Fourth, the book is commendably sceptical without being uniformly anti-technology. Its argument does not require the conclusion that AI should be prohibited. Instead, it asks whether a particular use is supported by evidence and whether the expected benefits justify the risks. This is compatible with a rights-based regulatory model. The Council of Europe's Framework Convention on AI and Human Rights, Democracy and the Rule of Law similarly adopts principles of human dignity, autonomy, equality, privacy, transparency, accountability and safe innovation, while requiring risk and impact assessment and providing for procedural safeguards and remedies.⁹

Finally, the book is highly effective as an interdisciplinary introduction. Lawyers need not become machine-learning engineers, but regulation requires enough technical literacy to distinguish genuine capability from exaggerated claims. Its pedagogical value lies in giving non-specialists a vocabulary for questioning technical assertions.

CRITICAL APPRAISAL: LIMITATIONS AND UNANSWERED LEGAL QUESTIONS

The book's principal limitation is the relative absence of a systematic legal architecture. It diagnoses why unreliable AI can be harmful, but it does not fully resolve who should bear responsibility when such systems are deployed. A rights-based framework requires answers to questions of standing, causation, standard of review, disclosure, evidentiary burdens, remedies and institutional competence. If a person is denied employment, insurance or a public benefit because of an algorithmic assessment, the legal system must determine not only whether the system was inaccurate but also what legal duty was

⁸ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference* 227–290 (Princeton University Press, Princeton, 2024).

⁹ Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature 5 September 2024.

breached and against whom the remedy lies.

A second limitation is that the book's treatment of discrimination could be developed further through comparative legal analysis. Existing scholarship shows that algorithmic systems may create difficult questions of disparate treatment and disparate impact, while appropriate documentation can sometimes make discriminatory decision-making easier to detect.¹⁰ This creates an important tension: opacity is a serious problem, but algorithmic decision-making can also produce records and repeatable processes that are more auditable than informal human discretion. A stronger legal treatment would distinguish between opacity of the technical model, secrecy of the institution and the practical ability of an affected person to challenge a decision.

Third, the book could say more about the global distribution of AI harms. Its institutional examples are particularly useful, but the legal and social consequences of AI differ across jurisdictions. In India, for instance, questions of privacy and automated data processing must be situated within the constitutional recognition of privacy and the statutory framework created by the Digital Personal Data Protection Act, 2023.¹¹ The DPDP framework regulates digital personal data processing, but it is not a comprehensive AI statute. This leaves important questions concerning high-impact automated decision-making, explainability, discrimination and public-sector AI outside a single dedicated regulatory framework.

The contemporary regulatory landscape also demonstrates how quickly the book's legal context can evolve. The EU AI Act is now moving from legislation to implementation: its main provisions apply from 2 August 2026, while prohibitions and AI-literacy obligations began applying earlier and certain governance and general-purpose AI provisions took effect in 2025.¹² The Council of Europe's Framework Convention, opened for signature in September 2024, represents another major development because it is a legally binding international treaty specifically connecting AI governance with human rights, democracy and the rule of law.¹³ These developments do not invalidate the authors' argument; rather, they expose the next question the book leaves underdeveloped: how should its empirical diagnosis be translated into

¹⁰ Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan and Cass R. Sunstein, "Discrimination in the Age of Algorithms", 10 *Journal of Legal Analysis* 113, 113–174 (2018).

¹¹ The Digital Personal Data Protection Act, 2023 (Act 22 of 2023).

¹² Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Art. 113.

¹³ Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature 5 September 2024.

enforceable legal obligations?

A further limitation concerns rapidly improving generative AI. Developments after the book's 2024 publication make some claims sensitive to technological change. This highlights a methodological problem for AI scholarship: an argument may remain sound while becoming incomplete as capabilities and regulation change. The 2025 paperback edition's new preface and epilogue itself illustrates this speed.¹⁴

Finally, the book could develop its political-economy analysis further. The authors recognise the danger of AI remaining controlled by largely unaccountable technology companies, but market concentration, control over data and computing capacity, intellectual property, labour displacement and platform governance deserve fuller treatment.

CONTEMPORARY RELEVANCE: FROM DIAGNOSIS TO RIGHTS-BASED GOVERNANCE

The book is especially valuable when read against contemporary regulatory developments. The EU AI Act illustrates the movement from broad ethical principles towards differentiated legal obligations based on risk. Its architecture recognises that not every AI application presents the same threat and that some practices warrant prohibition, while others require additional safeguards.¹⁵ This approach reinforces the authors' insistence that "AI" should not be treated as one homogeneous technology.

The Council of Europe Framework Convention goes further in a direction particularly relevant to the *International Journal of Human Rights Law Review*. It requires attention to human dignity and autonomy, equality and non-discrimination, privacy and data protection, transparency and oversight, accountability and responsibility, reliability and safe innovation. It also contemplates procedural safeguards, the ability of affected persons to challenge significant decisions, complaints mechanisms and iterative risk and impact assessments.¹⁶ These principles translate many of the concerns identified by Narayanan and Kapoor into a legal vocabulary.

India presents a different but equally important setting. The Digital

¹⁴ Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference* (paperback edn., Princeton University Press, Princeton, 2025).

¹⁵ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024.

¹⁶ Council of Europe, *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law* (CETS No. 225), 2024.

Personal Data Protection Act, 2023 recognises both the right of individuals to protect personal data and the need to process such data for lawful purposes.¹⁷ The Digital Personal Data Protection Rules, 2025 were subsequently notified, demonstrating that India's data-governance architecture is moving from statutory enactment towards implementation.¹⁸ Yet data protection alone cannot answer every AI governance question. An AI system can cause harm through discriminatory inference, unreliable classification or denial of opportunity even where the underlying processing is formally lawful. The book therefore helps reveal the limits of a purely data-centric regulatory approach.

OVERALL SCHOLARLY CONTRIBUTION

AI Snake Oil makes a significant contribution by challenging two misleading extremes in AI discourse: technological utopianism and indiscriminate technological pessimism. Its more useful position is that AI must be evaluated according to the task, evidence and institutional setting. That approach is intellectually modest but legally powerful. Law is concerned with consequences, responsibility and justification, and each of those questions requires a realistic understanding of what technology can actually accomplish.

The book's contribution is therefore diagnostic rather than regulatory. It does not provide a complete theory of AI governance, but it exposes assumptions that such a theory must confront. A legal framework based only on disclosure or human oversight may fail if institutions rely on systems fundamentally incapable of producing reliable predictions. A right to an explanation has limited value if the underlying decision should never have been automated.

For legal academics and policymakers, the book suggests shifting the regulatory focus from technology in the abstract to its institutional uses. The relevant regulatory unit should often be the AI-assisted decision and its consequences rather than "AI" as a technological category.

CONCLUSION

AI Snake Oil is an important intervention in a field where technological enthusiasm frequently moves faster than institutional understanding. Narayanan and Kapoor succeed in showing that the most immediate danger from AI is not necessarily a hypothetical superintelligence, but

¹⁷ The Digital Personal Data Protection Act, 2023 (Act 22 of 2023).

¹⁸ Ministry of Electronics and Information Technology, Government of India, Digital Personal Data Protection Rules, 2025, notified 14 November 2025.

the ordinary deployment of unreliable systems by institutions with real power over people. Their distinction between different forms of AI, insistence on empirical scrutiny and analysis of institutional incentives provide a valuable foundation for legal scholarship.

Its limitations are equally instructive. A legal reader will want a more developed account of rights, remedies, liability, discrimination, procedural fairness and regulatory enforcement. The rapidly evolving landscape of the EU AI Act, the Council of Europe Framework Convention and India's emerging data-governance framework further demonstrates that the next stage of scholarship must connect technical evidence with enforceable legal standards.

Ultimately, the book changes the regulatory question from "How should we regulate AI?" to a more precise and demanding question: "Under what conditions should a particular AI system be permitted to influence a particular human decision?" That reframing is its central scholarly achievement. It places institutional responsibility back at the centre of the debate and rejects the idea that technological complexity can itself justify deference. For human-rights law, this is a particularly valuable lesson. Technology may alter the means through which power is exercised, but it does not remove the legal obligation to justify that power, protect individual rights and provide meaningful remedies when automated systems fail.