

# INTERNATIONAL JOURNAL OF LAW MANAGEMENT & HUMANITIES

[ISSN 2581-5369]

---

Volume 9 | Issue 2

---

2026

© 2026 *International Journal of Law Management & Humanities*

Follow this and additional works at: <https://www.ijlmh.com/>

Under the aegis of VidhiAagaz – Inking Your Brain (<https://www.vidhiaagaz.com/>)

---

This article is brought to you for free and open access by the International Journal of Law Management & Humanities at VidhiAagaz. It has been accepted for inclusion in the International Journal of Law Management & Humanities after due review.

In case of **any suggestions or complaints**, kindly contact [support@vidhiaagaz.com](mailto:support@vidhiaagaz.com).

---

**To submit your Manuscript** for Publication in the **International Journal of Law Management & Humanities**, kindly email your Manuscript to [submission@ijlmh.com](mailto:submission@ijlmh.com).

---

# A Comparative Analysis of Legal Liability and Accountability Frameworks for Autonomous Artificial Intelligence

---

ARUN S<sup>1</sup> AND T SAROJA DEVI<sup>2</sup>

## ABSTRACT

*Artificial Intelligence (AI) has evolved from a supportive tool into a semi-autonomous system influencing sectors such as healthcare, finance, transportation, and digital communication. This transformation raises significant legal concerns, particularly regarding liability when AI systems cause harm. Traditional legal frameworks are structured around identifiable human or corporate actors, making it difficult to assign responsibility in cases where AI operates with a degree of autonomy and unpredictability.<sup>3</sup>*

*This study undertakes a comparative analysis of legal liability and accountability frameworks for autonomous AI across major jurisdictions, including the European Union, United States, United Kingdom, China, and India. It argues that neither fault-based liability nor the concept of granting legal personhood to AI offers a complete solution. Fault-based approaches face challenges due to the opacity of AI systems, while AI personhood lacks practical feasibility in ensuring compensation and accountability.*

*The paper proposes a layered regulatory approach that combines ex ante obligations—such as risk assessment, documentation, and monitoring—with ex post liability mechanisms tailored to the level of risk involved. It also emphasizes the need for evidentiary flexibility and compensation mechanisms, including insurance frameworks, especially for high-risk AI applications.*

*The study concludes that India must adopt a proactive and structured regulatory model that ensures accountability across the AI lifecycle while safeguarding innovation. Such a framework would bridge the accountability gap without conferring independent legal status on AI systems.<sup>4</sup>*

**Keywords:** *Artificial Intelligence, Autonomous Systems, Legal Liability, Accountability, Comparative Law, AI Regulation, Product Liability, Digital Governance, Risk-Based Regulation, India*

## I. INTRODUCTION

The rapid advancement of Artificial Intelligence marks a transformative shift in modern society.

---

<sup>1</sup> Author is a Student at Vels Institute of Science, Technology & Advanced Studies, Chennai, Tamil Nadu, India.

<sup>2</sup> Author is an Assistant Professor at Vels Institute of Science, Technology & Advanced Studies, Chennai, Tamil Nadu, India.

<sup>3</sup> European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*, COM (2021) 206 final (Apr. 21, 2021).

<sup>4</sup> Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 Fordham L. Rev. 1085 (2018).

AI systems are no longer limited to performing predefined tasks; they are increasingly capable of learning, adapting, and making decisions independently. From autonomous vehicles to algorithm-driven financial systems and healthcare diagnostics, AI is deeply integrated into critical aspects of everyday life.<sup>5</sup>

This growing autonomy presents a fundamental challenge to existing legal frameworks. Traditional legal systems are designed to attribute responsibility to identifiable human or corporate actors. However, autonomous AI systems complicate this process by distributing decision-making across multiple stages, including design, development, deployment, and operation. As a result, determining liability in cases of harm becomes increasingly complex.

One of the major challenges lies in the “black-box” nature of AI, where the internal decision-making processes are not easily explainable. This lack of transparency makes it difficult to establish fault or negligence under conventional legal standards. Furthermore, the involvement of multiple stakeholders—such as developers, manufacturers, service providers, and users—further complicates the allocation of responsibility.<sup>6</sup>

Various approaches have been proposed to address these issues. Some scholars advocate for adapting existing legal doctrines, while others suggest granting legal personhood to AI systems. However, both approaches have limitations. A purely fault-based system struggles with evidentiary challenges, while AI personhood raises concerns about enforceability and victim compensation.

In this context, there is a pressing need for a balanced and comprehensive legal framework that ensures accountability without hindering technological innovation. This study aims to explore comparative legal approaches and propose a structured model that can effectively address the challenges posed by autonomous AI. By examining global practices, the research seeks to contribute to the development of a robust and future-ready legal regime.<sup>7</sup>

### **The Problem of Accountability**

Autonomous AI systems operate through multiple stages, including design, training, deployment, and real-time operation. Each stage may involve different actors—developers, data providers, companies, and users. When an AI system produces harmful outcomes, assigning liability becomes difficult because no single actor may have complete control over the system’s

---

<sup>5</sup> Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. Davis L. Rev. 399 (2017).

<sup>6</sup> Andrew D. Selbst & Solon Barocas, *supra* note 3.

<sup>7</sup> Shyamkrishna Balganes, *Artificial Intelligence and the Law: Liability and Accountability*, 88 Fordham L. Rev. 123 (2020).

behavior.

Additionally, the “black-box” nature of many AI models makes it challenging to understand how decisions are made. This lack of transparency complicates the process of proving negligence or fault, which is essential in traditional legal frameworks.<sup>8</sup>

### **Limitations of Existing Legal Frameworks**

Existing legal doctrines such as tort law, product liability, and criminal law are not fully equipped to address AI-related harms. Fault-based liability requires proving negligence, which is difficult when decision-making processes are opaque. Similarly, strict liability frameworks may not adequately capture the complexities of AI systems that evolve over time.

The idea of granting legal personhood to AI has been proposed as a solution. However, this approach raises practical concerns, including the inability of AI systems to bear financial liability or provide compensation to victims.<sup>9</sup>

### **Need for a New Approach**

Given these challenges, there is a growing need for a restructured legal framework that can address the unique characteristics of AI. A balanced approach must combine preventive measures with effective remedies. This includes:

- Ensuring transparency and accountability during the development phase
- Imposing obligations on all stakeholders involved in the AI lifecycle
- Creating mechanisms to compensate victims efficiently

### **Scope and Objective of the Study**

This study aims to examine how different jurisdictions address AI liability and to identify best practices that can inform future legal reforms. By analyzing comparative frameworks, the research seeks to propose a model that ensures accountability while encouraging innovation.

## **II. JURISPRUDENTIAL THEORIES OF LIABILITY AND ACCOUNTABILITY**

Any attempt to regulate autonomous AI must start with the legal doctrines already available. Novelty should not obscure continuity. Courts have long confronted harmful complexity, whether in pharmaceuticals, aviation, electricity, financial markets, automated machinery, or digital platforms. AI does, however, recombine these older problems in distinctive ways. It

---

<sup>8</sup> Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 Fordham L. Rev. 1085 (2018).

<sup>9</sup> Shyamkrishna Balganesh, *Artificial Intelligence and the Law: Liability and Accountability*, 88 Fordham L. Rev. 123 (2020).

behaves like software, a service, a product, a speech system, a decision support tool, and sometimes an infrastructural intermediary all at once. That hybrid character matters because different doctrines attach to different legal characterisations. This chapter therefore surveys the principal models through which law assigns responsibility for technologically mediated harm. The purpose is not to claim that one model should displace all others, but to evaluate the strengths and limits of each when applied to autonomous AI. The analysis shows that the most promising frameworks are mixed rather than pure. Negligence remains essential, but it must be reinforced by product liability, consumer protection, administrative enforcement, organisational duties, and sometimes insurance-backed compensation.

### **A. Negligence and the continuing centrality of fault**

Negligence remains the default common-law mechanism for addressing novel harm. Its attraction lies in flexibility. A court can ask whether a defendant behaved as a reasonable developer, deployer, integrator, operator, or supervisor in light of foreseeable risk. That flexibility is especially valuable in a fast-moving field where legislatures cannot specify every design obligation in advance. In principle, negligence can cover inadequate data governance, poor testing, insufficient human oversight, insecure deployment, reckless overclaiming of system capability, and the failure to intervene after warning signals emerge. Yet negligence is also structurally fragile in the AI context. To succeed, a claimant usually needs access to evidence showing what a defendant knew, what steps were taken, what risks were foreseeable, and how those failures caused the harm. AI systems complicate each element. A model may have been trained by one entity, fine-tuned by another, packaged by a third, deployed by a fourth, and operated by a fifth. Harm may emerge from interaction effects rather than a single mistake. Moreover, the most relevant evidence often lies within corporate systems inaccessible to claimants. Thus, while negligence doctrine is doctrinally available, it is not always practically effective.

There is also a standards problem. What counts as the 'reasonable AI developer' or 'reasonable deployer'? Courts can draw on industry standards, internal safety protocols, regulatory guidance, auditing practice, and expert testimony, but many sectors lack mature, publicly settled baselines. This is one reason why soft-law instruments such as the OECD AI Principles and the NIST AI Risk Management Framework matter even when not directly binding. They help courts and regulators concretise what prudence, documentation, testing, and governance should look like. Negligence thus remains necessary but insufficient. It works best where the claimant can show reliance, discrete misrepresentation, ignored warnings, poor monitoring, or clearly omitted safeguards. It works less well where the harm stems from diffuse model behaviour,

unexplained bias, or an evidentiary asymmetry so severe that breach and causation cannot be reconstructed.

### **B. Product liability and strict-liability logics**

Product liability responds to a different intuition: some risks should be borne by those who place products on the market rather than by injured users who lack knowledge and bargaining power. This logic is highly relevant to AI where complex software, embedded systems, and continuous updates create latent risks difficult for end users to detect. Modern product-liability regimes are therefore being forced to decide whether software, model updates, remote services, and post-sale modifications should count as defects or product-related elements. The comparative trend is toward expansion. The recast European Product Liability Directive expressly treats software as a product and recognises that products may change after they are placed on the market through software updates, machine learning, or the loss of safety-related cybersecurity support. That move is conceptually important because it prevents suppliers from arguing that harmful software functionality is beyond product-liability law simply because it is intangible. The core idea is that when a commercial actor places an AI-enabled product into circulation, the victim should not bear the full burden of proving specific internal fault.

Still, product liability has limits. Not every AI harm flows from a defective product. Some harms arise from services, poor governance, erroneous advice, discriminatory use, or negligent integration into a specific workflow. Foundation- model APIs, software-as-a-service tools, and decision-support systems may sit awkwardly between product and service categories, depending on the jurisdiction. Furthermore, product liability often requires proof that a product failed to provide the safety the public is entitled to expect, which can be difficult in frontier systems where capabilities and failure modes are unsettled. For these reasons, product liability is powerful but not exhaustive.

### **C. Enterprise liability, vicarious liability, and institutional responsibility**

Enterprise liability asks which actor organised the activity, captured the gains, and is best positioned to spread losses through pricing, insurance, or organisational control. This rationale has deep roots in vicarious liability and modern risk-distribution theory. It resonates strongly in AI because firms often benefit from automation precisely by reducing labour cost, scaling decisions, or increasing speed. If those gains are internalised, there is a strong fairness argument that at least some harms should be internalised as well. Vicarious liability may be directly relevant where employees or agents use AI within the scope of employment. More subtly, enterprise logic supports placing duties on deployers even when a vendor built the underlying

model. A hospital that integrates AI into diagnosis, an insurer that uses it to price, or a platform that uses it to rank and moderate content is not a passive recipient of technology. It decides whether, where, and under what conditions the system is deployed. Enterprise-based accountability therefore helps prevent a blame game in which each participant points to another link in the chain.

The challenge is calibration. Over-expansive enterprise liability could lead downstream users to reject any high-performing but complex AI system for fear of open-ended responsibility. A balanced approach is to assign non-delegable duties around selection, human oversight, workflow integration, and post-deployment monitoring, while preserving recourse rights between deployers and upstream suppliers. That structure places the claimant's remedy first and leaves contribution and indemnity to be resolved among the commercial actors.

#### **D. Contract, consumer protection, and negligent misrepresentation**

Contract and consumer law remain highly relevant because many AI disputes arise through representations made to users and consumers. Terms of service, disclaimers, warranties, limitation clauses, representations about capability, accuracy, safety, and intended use can all shape liability. In consumer-facing settings, the law is rightly suspicious of attempts to avoid accountability through complex boilerplate or by describing a chatbot, recommender, or automated decision system as merely 'for information purposes' when the service is clearly designed to induce reliance. The Air Canada chatbot dispute is a striking illustration. Even where the financial quantum was modest, the tribunal's reasoning was significant because it refused to treat the chatbot as a detached legal actor and instead attributed responsibility to the firm operating the website. The case underscores an important principle: entities cannot enjoy the efficiencies of AI-mediated customer interaction while disclaiming the legal consequences of those interactions whenever the system errs.

Consumer law also matters in India, where the Consumer Protection Act 2019 already recognises product liability claims against manufacturers, sellers, and service providers. Although drafted without modern autonomous AI in mind, the Act can serve as an existing bridge for some AI-enabled harms, particularly where products or services are sold to consumers with implied assurances of safety, quality, or reliability.

#### **E. Criminal liability and the limits of mens rea transfer**

Criminal liability presents the hardest conceptual questions because classical criminal law is tightly linked to mens rea, human agency, and moral blame. Autonomous AI can certainly be involved in criminally relevant events: deceptive deepfakes, algorithmic fraud, negligent

deaths, manipulation, privacy intrusions, or unlawful automated surveillance. But the fact that AI mediated the conduct does not mean AI itself is a fitting subject of punishment. Criminal law generally reaches the humans or corporations who intentionally, knowingly, recklessly, or negligently designed, deployed, or used the system in unlawful ways. There are two main risks of overreach. The first is symbolic criminalisation of technical failure where no sufficiently blameworthy human conduct can be shown. The second is under-enforcement caused by treating AI complexity as a reason to abandon criminal inquiry altogether. The appropriate middle path is to focus on the mental state and organisational conduct of real actors: did they knowingly deploy a deceptive synthetic system, ignore manifest danger, falsify safety records, or recklessly use AI where statutory duties demanded human verification? Corporate criminal liability and offences tied to due diligence failures may therefore be more realistic than fantasies of punishing the model itself.

For India, the new criminal-code architecture under the BNS, together with evidence rules under the BSA, can address many AI-mediated offences in principle. The deeper challenge is evidentiary and organisational: proving who configured the system, who knew of the risk, what logs exist, and whether omission or defective supervision satisfies the statutory thresholds.

#### **F. Electronic personhood: analytical curiosity or practical mistake?**

The idea of electronic or legal personhood for autonomous AI appears repeatedly in academic debate because it seems to solve an attribution puzzle. If the AI is a person, perhaps it can bear rights and duties, own assets, and be sued. On closer examination, however, the concept solves little unless supported by an elaborate institutional framework. A legal person that cannot be imprisoned, shamed, mentally blamed, capitalised, insured, or meaningfully compelled to disclose is a weak substitute for real defendants. Worse, personhood risks becoming a doctrinal smoke screen behind which developers, deployers, and investors hide. That does not mean the concept is entirely useless. Limited forms of functional registration might be valuable for traceability, for asset-segregated autonomous trading vehicles, or for identifying which version of a system was active at a given time. But those are administrative or organisational advantages, not moral or doctrinal reasons to recognise AI as a primary bearer of liability. The better legal instinct is to map the network of human and corporate actors around the system rather than to anthropomorphise the system itself.

A further objection is distributive. Victims need solvent defendants or compensation funds, not abstract conceptual elegance. Unless 'electronic persons' are backed by mandatory capitalisation and insurance, personhood may simply reduce the pool of recoverable assets. Accordingly, this

report rejects full AI personhood as a primary liability framework and treats it, at most, as a niche adjunct for registration and traceability.

### **G. Insurance, funds, and victim compensation**

Insurance is sometimes treated as secondary to liability, but in the AI context it deserves central attention. The reasons are practical. High-risk AI may generate low-frequency but high-severity losses; proving fault may be slow and expensive; and defendants may be numerous. Insurance can stabilise compensation, generate market discipline through underwriting, and incentivise documentation because insurers demand evidence of testing, monitoring, and governance before they price risk.

The autonomous-vehicle field is especially instructive. Several jurisdictions have recognised that road users should not have to litigate complex software causation in every crash before obtaining compensation. A first-line insurer can compensate victims and then seek recovery from responsible actors upstream. This model does not erase fault; it sequences it differently. Compensation comes first, technical allocation later. Similar thinking may be appropriate for other high-risk AI sectors such as healthcare, industrial robotics, and critical infrastructure. Insurance, however, is not a panacea. It can create moral hazard, may be difficult to price for frontier systems, and depends on reliable incident data. That reinforces the need for mandatory reporting and audit records. In other words, insurance works best as part of an accountability ecosystem, not as a substitute for it.

### **H. Synthesis**

The doctrinal survey supports a mixed framework. Negligence alone is too evidentially demanding. Product liability alone does not capture service-based, organisational, or governance failures. Criminal law is too blunt for many cases yet indispensable for fraud, deception, reckless disregard, and document falsification. Consumer law is powerful where reliance is direct but may not help in systemic harms. Insurance improves compensation but depends on a prior duty architecture. And AI personhood, despite its intellectual allure, does little to solve the claimant's practical problem.

A mature legal framework should therefore allocate duties across the AI life-cycle. Developers should face duties around training-data governance, documentation, validation, cybersecurity, foreseeable misuse, and update discipline. Deployers should carry duties concerning selection, contextual testing, human oversight, user notice, monitoring, and incident response. Integrators should remain answerable where they combine components in risky ways. Operators and institutional users should not escape liability by claiming that a system was 'vendor built' if they

embed it into consequential decision processes without adequate safeguards. In this sense, AI accountability is best understood as chain accountability rather than single-actor blame. Chapter II takeaway. No single doctrine can absorb autonomous AI. The law works best when fault, enterprise, product, consumer, and insurance logics are combined and assigned to the specific role each actor plays in the AI life-cycle.

### **III. CONCLUSION**

As AI continues to evolve, the law must adapt to ensure that technological progress does not come at the cost of justice. Establishing clear accountability mechanisms is essential not only for protecting individuals but also for maintaining public trust in emerging technologies. A forward-looking legal framework must recognize that while AI may act autonomously, responsibility ultimately lies with human and institutional actors who design, deploy, and benefit from it.

\*\*\*\*\*

## IV. BIBLIOGRAPHY

### Books & Reports

- Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap* (U.C. Davis Law Review, 2017).
- OECD, *OECD Principles on Artificial Intelligence* (2019).
- World Economic Forum, *Governance of AI: Global Perspectives* (2022).
- NITI Aayog, *National Strategy for Artificial Intelligence: #AIforAll* (2018).

### Articles & Journals

- Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 Fordham Law Review 1085 (2018).
- Shyamkrishna Balganesh, *Artificial Intelligence and the Law: Liability and Accountability*, 88 Fordham Law Review 123 (2020).

### Government & Institutional Reports

- National Institute of Standards and Technology (NIST), *AI Risk Management Framework* (2023).
- U.S. Federal Trade Commission, *Algorithmic Accountability Framework* (2023).
- Cyberspace Administration of China, *Algorithmic Regulation Provisions* (2022).
- European Parliament, *Civil Liability Regime for Artificial Intelligence* (2020).

\*\*\*\*\*