Journal of Intellectual Property Rights
Vol 31, July 2026, pp 523-534
DOI: 10.56042/jipr.v31i4.18015

Enterprise Liability and Risk Governance in the Regulation of Artificial Intelligence

Aditi Bharti & Gagandeep Kaur†
School of Law, UPES, Dehradun – 248 007, India

Received: 18th March 2025; revised 5th January 2026

The growing autonomy of artificial intelligence systems is reshaping long-standing assumptions about legal responsibility. Traditional tort frameworks were built around human actors and clear causal links. These assumptions weaken when harm is produced by systems that learn, adapt and operate with limited human oversight. The paper examines how AI challenges established legal reasoning by analysing liability through two key dimensions. The first concerns the responsibility of human actors like developers, deployers and professional users, whose decisions shape AI behaviour. The second considers whether, and in what sense, liability can be attributed to AI systems themselves. The analysis also draws on concepts like enterprise liability, Ulrich Beck's risk society theory and European Union's risk based approach to AI regulation. The paper argues for a combined model that uses regulatory duties to prevent harm and tort law to address such harm. This framework aims to balance accountability, fairness and innovation.

Keywords: Artificial Intelligence, Enterprise Liability, Legal Personhood, Risk Society, Tort Law

In March 2018, the streets of Tempe, Arizona, became the setting for a defining legal and ethical question of the twenty-first century. An autonomous Uber test vehicle struck and killed a pedestrian, Elaine Herzberg, while operating in self-driving mode. The vehicle's sensors detected the victim, but the system failed to classify her as a collision risk in time to act. The safety driver, who had momentarily looked away, was later charged with negligent homicide, while Uber faced no criminal liability.[1] This incident exposed a profound gap in the existing legal framework: when artificial intelligence acts autonomously, who should be held responsible the human operator, the company, or the algorithm itself? The case illustrates how traditional principles of tort law, grounded in human agency and intent, are increasingly strained by the rise of autonomous and intelligent technologies. It also raises a deeper jurisprudential question: can the law of liability, designed for human action, meaningfully govern artificial agents whose operations transcend human foresight and control? Against this backdrop, this paper examines how AI interacts with legal reasoning processes, exploring its implications for liability, responsibility, and regulation. By situating the analysis within tort law and the risk-based governance framework emerging in the European Union, the study seeks to uncover how legal systems can adapt to the complexities of technological autonomy while preserving the normative foundation of human accountability.

Artificial intelligence is often referred to as the fourth *industrial revolution* due to its rapid advancement and the extent of its impact on society. In recent years, governments worldwide have been facing a challenging task: how to regulate artificial intelligence because this technology has now started to impact humans' daily lives.

The unique feature of artificial intelligence is its autonomy. Machine learning and self-driving vehicles are prime examples of artificial intelligence that operate without human assistance. These systems can reason, plan, and communicate to provide valuable assistance to human beings. However, even though these technologies are beneficial, they can also cause significant harm to humans. The existing regulatory framework is ill-equipped to govern and regulate the complexities of artificial intelligence, majorly because the law is not evolving at the pace of the technology. The legislature can be considered partially responsible as the law cannot keep pace with the growing technology since artificial intelligence is still evolving, and it needs to know how or what output it will generate or what response it will give, thereby making it uncertain. Legal and ethical regulations, on the other hand, require certainty to regulate the

†Corresponding author: Email: gkaur@ddn.upes.ac.in

subject matter, and the rate at which AI is developing makes it challenging to establish appropriate rules, highlighting the need to balance innovation with protection.

Although there is no standard definition of artificial intelligence, various attempts have been made to define it. Dr. Ahuja, in his paper cites the definition of artificial intelligence by Philips Jakson, “the science of making computers do things that people would say requires intelligence”. He also cites Ray Kurzweil who defined artificial intelligence as “the science of making computers do things that require intelligence when done by humans”.[2] Similarly, Zohuri and Behgounia in their work describes artificial intelligence as “the ability of a computer program or a machine to think and learn, similar to the way that human does naturally and can be as cognitive as possible via it’s artificial neural network built in the machine”.[3] Even though there is a lack of consensus on the definition, every definition agrees on one issue that this technology can mimic or function similar to human intelligence, like thinking, analysing, and learning from experience. As the name suggests, these different entities function according to their level of intelligence. Even Though there has been a lot of discussion around regulating AIs, the level of Artificial General Intelligence and Artificial Super Intelligence is yet to be achieved by the technology. Nevertheless, looking at the pace of development, these technologies would soon become reality and an integral part of our lives.

The fundamental question to determine the responsibility of artificial intelligence would be raised due to its advancements. The issue would revolve around fixing the responsibility of a human element, like the user or manufacturer of the AI. These issues require reconsidering the suitability of the existing legal framework.

This paper explores the evolving relationship between artificial intelligence and the legal reasoning process, focusing on how AI challenges, complements, and transforms traditional frameworks of legal responsibility. It examines the intersection between technological autonomy and legal accountability, situating these debates within the broader context of tort law and risk regulation. The central inquiry concerns how law can respond to the unique risks created by AI, risks that are diffuse, systemic, and often untraceable to a single human actor. To address this, the paper is divided into two main parts. The first part analyses human accountability in the operation and deployment of weak AI systems, considering how doctrines such as negligence, vicarious liability, and enterprise liability can adapt to new forms of technological mediation. The second part investigates whether, and to what extent, legal liability can be attributed to AI systems themselves, engaging with the ongoing debate on AI personhood and the conceptual boundaries of legal subjectivity. Integrating Ulrich Beck’s *risk society theory* and the European Union’s risk-based approach to AI governance, the paper argues for a hybrid model that combines preventive regulation with responsive tort principles. Through this approach, it seeks to demonstrate how law can evolve to manage the “manufactured risks” of AI while maintaining the moral and legal coherence of human responsibility.

Analysing Rights and Liability of a *Person*

Any discussion on rights and liability involves an understanding of the legal relationship between person and entity along with the surrounding concepts like entitlements, obligations, and consequences of the actions. According to Salmond, legal rights are "*interest that is protected and acknowledged by the rule of law"* emphasising that there is a duty associated with this interest, and it is morally wrong to disregard it.[4] Similarly, for any person to be held liable under criminal law, two elements are required: *actus reus* and *mens rea*. The subjective element, i.e., mens rea, requires culpability, whether intentional or not. Intentional culpability means that the person can understand and predict the outcome of his action. Negligence involves carelessness and some level of awareness of the potential consequences. Additionally,the act should result in infringement.[5] Meanwhile, in civil law cases, liability arises from tort, which includes infringement of other's rights.

These fundamental concepts get complicated when applied to artificial intelligence entities. For any right or liability to be assigned, the entity should be recognized by law as a person, either natural or juristic, neither of which has been ascribed to AI. However, recognizing artificial intelligence as a juristic person would raise several issues regarding allocating its rights and duties. Like under civil law, an AI to be held liable for its acts would need a juristic personality, similar to how law holds corporations criminally liable for their actions. Considering corporate criminal liability as an example, corporations are also considered legal fiction and are held responsible for the actions of their employees.[6] In

contrast, an AI would be held accountable for its actions, not attributed to any individual. However, this does not need any explanation that holding AI liable for its actions will not be this straightforward.

The Black-Box Dilemma

Artificial intelligence relies heavily on problem solving algorithms to perform the tasks assigned not only effectively but also quickly. However, understanding how AI functions and processes the information fed to it is often an impossible task for humans. This is an important issue especially with deep learning models and neural networks. Artificial intelligence functions on complex algorithms developed by the software to attain the desired result wherein the neural network divides the problem into multiple pieces and process them to generate the output. In this process, the algorithm is not written by a human, instead by the software itself. Since there is no access to know how the neural network is functioning and processing the data, it is termed as *black-box issue.*[7]

One of the significant challenges posed by the black box is the issue of determining accountability and responsibility. A situation where the AI-generated decision affects rights and interests, determining liability becomes impossible, given the opaque nature of the decision-making process. The opacity of black box hinders the process of determining liability and fixing responsibility. However, it would be too soon to determine if the black-box nature of AI would pose legal problems in the future or not in the decision-making process. Lack of knowledge about the decision-making process of AI would undermine the ability to demonstrate illegal action or omission.[8]

Considering Artificial Intelligence as an Artificial Person

Conferring rights and liabilities to artificial intelligence depends on its status as a legal person. However, granting the status of legal person is a complex process which is dependent not only on the legislature but also on the public policy. While the concept may seem revolutionary, like it did when corporations were recognised as legal persons, a thorough examination must be done before granting the status of personhood to artificial intelligence while ensuring that they can be held responsible for their actions. The personality theory given by Kelson asserts that the status of legal personality on a non-human entity is the attribution of a human character to the entity in order to establish the rights and duties on it.[9] Similarly, in the *Hohfeldian analysis* of rights and obligations, they are considered as jural correlative of each other. Therefore, when the issue of granting rights and liabilities to artificial intelligence is discussed, the question is raised on its effectiveness. Recognizing an artificial intelligence as a juristic person could also reduce the liability of humans involved in either developing or using the technology.[10] However, the real challenge is not whether AI should have a legal personality or not; it is in deciding the appropriate model of personhood for AI. Granting legal recognition to AI's personality would create civil liability, typically addressed through monetary compensation. This would create a legal separation between AI and its owner, similar to corporate personality, where the assets of a company are separate from those of its owner. Unlike a corporation, AI does not possess its own assets, and if AI were to pay compensation, it would lack the required resources. Therefore, to make good any loss caused by AI, it needs to have assets. Contrary to this, in case of criminal liability, vicarious liability is often employed holding the company management responsible for the actions of the corporation. Similar concepts can be applied to AI. However, given the autonomous nature of AI wherein it can take its own decision, it would be unjust to punish the human counterpart for the actions of AI. The nature of AI would make it unfair to convict a human under conventional legal framework. The legislature will have to come up with other punitive measures to hold artificial intelligence accountable like, deactivation or revocation of operation license or personhood status etc. Deciding upon this would require a detailed analysis while exploring relevant types of punishment or penalty models that could help legislation achieve its purpose.[11]

Although this thought process is too hypothetical, it might be helpful in determining the circumstances under which an artificial intelligence can hold rights and liabilities. The concept of *humans* and the prerequisite of becoming a *person* must be discussed to address this. As discussed earlier, humans are endowed with the status of personhood by the virtue of being born and thus are granted rights and duties by law. Fauzan draws an interesting analysis of how law and morality are intertwined. He contends that the legal concepts of personhood, along with the associated rights and duties, are derived from moral concepts. To be considered a legal person, one must be an ethical subject first. Fauzan takes the help of Hobbes, who states that since no morality exists in the natural state, humans, who are selfish by nature, are free to act freely. It is because of

this selfishness that moral values were developed to ensure self-preservation and rationality. Hobbes states that humans have evolved over the years to live communally to survive, which has consequently promoted values like empathy and altruism, which underpins moral behaviour. Morality functions by creating rights and duties, leading to the concept of moral standing whereas rights make sense only when there is a duty to maintain societal relationships. Therefore, an entity must have moral standing to bear rights and obligations, meaning it must be entitled to ethical considerations. It is necessary to differentiate between ethical and legal rights and responsibilities. Legal rights and duties are created by law, and the criteria for recognizing humans as the only natural persons stem from their status as moral subjects. Fauzan concludes by stating that humans' ability to experience pain and pleasure and to behave in a survival-favorable manner through reward and punishment-based behavior makes them bear rights and duties towards society. [12]

Enterprise Liability v *Bundle Personhood*

Since the existing framework does not recognize AI as a person, attributing liability to artificial intelligence is a futuristic concept. Currently, machines are incapable of bearing duties, which is a fundamental aspect of the idea of liability. At the heart of this debate are the notions of accountability and liability, often misconstrued as synonymous. In this reference, two theoretical approaches have been analysed in this paper: one by Lauren Geisser which adapts the concept of enterprise liability[13] and the other by Visa Kurki, who developed the concept of bundle theory.[14] Geisser in her paper applies the concept of enterprise liability to hold Uber liable for the accident and subsequent death by their car. She advocates for accidents and harms caused by technology-driven business to be absorbed by the firm as a cost of doing business. Kurki, on the other hand, in his book treats personhood as a flexible bundle of legal incidents. It permits the entity to be a legal person for some purposes and others. Contrasting these approaches gives different visions for AI governance. One focuses on corporate accountability while other on reconfiguring the personhood itself.

Inspired by vicarious liability theory, Geisser in her paper advocates for enterprise liability for holding technology based companies liable for harm caused by them. In her analysis of Uber fatality case, she argues that Uber should be held vicariously liable for the driver caused accidents. Since ridesharing is an inherent business activity of the company, the liability should also lie with the company and not the driver. Therefore, the enterprise-based model applies conventional tort liability on AI-driven firms to manage and mitigate the risks. By contrast, Kurki in his *bundle theory of legal personhood* rejects the binary view of personhood, describing it as interconnected incidents. This allows an entity to be a "person" by holding different subsets of legal incidents without enjoying all the incidents of full personhood. This theory provides unprecedented flexibility when determining the personhood status of an AI entity. If an AI entity can satisfy criteria like agency or asset ownership, it could be granted a form of legal personhood. This theory eliminates any border between legal personhood and non-personhood, allowing an AI to be a legal person for certain liabilities and transactions but not others.

These theories highlights key stakes in the AI law debate. The enterprise-liability approach presumes that the responsibility should vest with the humans behind the AI. This liability model is generally favoured by the regulators. For example, the EU's AI Act explicitly avoids granting legal personality to AI. Earlier drafts of the Act recommended *electronic personality* for artificial intelligence, which were rejected as it might shield companies from accountability. Instead, the Act creates obligations for the human and corporate actors deploying AI. In contrast, the bundle theory suggests expanding the legal term itself. A creative personhood framework such as electronic personality might prove to be a more nuanced way to distribute liability. However, emerging solutions to AI harm will likely combine elements from both theories, for example, a strict liability model along with dependent personhood.

Tort Law Foundations

The common law of torts grounds liability primarily in duties to avoid harming others, enforced through negligence, strict liability, or vicarious liability doctrines. A negligent actor is liable only if a legal duty was breached, a causal link to injury is proven, and damages resulted.[15] Courts thus traditionally require a breach of duty by a human actor.[15] In medical and product contexts, for example, traditional tort causes of action include malpractice (physician negligence), respondeat superior (employer or hospital vicarious liability for employee practitioners), and products liability against manufacturers.[16]

Each of these schemes embodies a form of "fault" or breach-based liability, in that liability turns on a failure

to exercise reasonable care (in negligence), failure in institutional supervision, or the presence of a defect in a marketed product. Thus, conventional tort law distributes losses partly on the basis of which actors are deemed at fault or failed to discharge their duties. The Restatement (Third) of Torts emphasizes that negligence (fault liability) is the general rule, while strict liability is limited to special circumstances; nevertheless, even strict liability was characterized as liability for risks "characteristic" of an enterprise.[17]

Under vicarious liability (respondeat superior), employers or institutions may be liable for the wrongs of their agents, but such liability still rests on agency principles and institutional duty (e.g. negligent hiring or supervision).[18]

Notably, even under these doctrines an underlying individual breach is typically required: hospitals for instance may be held liable if they negligently credential or supervise physicians, but absent such negligence only the individual tortfeasor (the physician) would normally be sued. Products liability imposes strict responsibility for injuries caused by a defective product, but even their duties (to design safely, to warn) are owed to identifiable consumers or intermediaries. In sum, tort law's foundational orientation is corrective: it recognizes a legal wrong (breach of duty) and compels wrongdoers to compensate victims. As Chan emphasises, "if a patient cannot show that a …actor breached [a] duty, then no tort has been committed, regardless of the injury".[15]

Enterprise Liability: Concept and Theory

Enterprise liability is a collective theory of tort responsibility that shifts the focus from individual fault to the risks inherent in a firm, industry, or activity. In its classical articulation, enterprise liability holds that an entire business enterprise (or group of related enterprises) should bear the costs of injuries "characteristic" of the enterprise's operations, regardless of which individual actor specifically caused a given injury.[19,20]

The defining idea is that accident costs should be borne by the enterprise as a whole and shared among those who profit from or participate in the enterprise. For example, Keating explains that under enterprise liability "the costs of an injury should be shared by those who profit from the activity responsible for the injury," rather than dispersing losses onto innocent victims or unrelated parties.[17] Similarly, John Fabian Witt's article defines enterprise liability as the notion that firms or activities should pay for the harms they cause "because they are in the best position to avoid causing such injuries, and because they are better able to spread the costs associated with them".[19] Thus, enterprise liability effectively amalgamates multiple actors (employees, owners, affiliated entities) into a single "enterprise" and imposes joint responsibility for harms flowing from that enterprise's characteristic risks. Scholars have varied in emphasizing the normative basis for enterprise liability. While many traditional theorists approached enterprise liability on economic or utilitarian grounds (focusing on compensation and deterrence), Keating revived a fairness-based rationale. He argues that tort rules are fair if they advantage all affected parties ex ante, even the worst-off in an accident scenario. From this perspective, fairness calls for spreading the cost of harmful accidents across the beneficiaries of the risky activity rather than concentrating it on an individual victim. As Keating explains, "accident costs of risky activities should be spread among all those who benefit from the activity, ideally in proportion to their degree of benefit," making enterprise liability "prima facie fairer than negligence liability". [21]

In short, enterprise theory is grounded on two core propositions: first, that an activity should bear its own accident costs (not random victims), and second, that those costs should be divided among the members or participants of the enterprise that imposed the risk.[17] These principles are often couched in terms of commutative justice or distributive fairness, mirroring Kantian and social-contract notions of shared obligations.[22] Proponents thus view enterprise liability as a form of strict liability for an industry or institution: it holds enterprises accountable for harms flowing from their agency, even absent individual fault.[23] In this light, enterprise liability is sometimes called a modern embodiment of strict liability: the enterprise itself (the "doer" as Keating notes) bears loss for characteristic accidents "rather than [placing losses] on the basis of fault".[24] Practically, enterprise liability often overlaps with or encompasses doctrines like vicarious liability, public welfare strict liability, and compensation schemes (e.g. workers' compensation), since in each the risks of an activity are socialized or collectivized. The critical innovation of enterprise liability theory is extending this collectivization further to hold entire institutions or industries accountable for harms that flow from the systemic risks they create, whether or not any individual was negligent.

Cases and Limitations

The doctrine of enterprise liability first emerged prominently in hospital malpractice law. Traditionally, hospitals escaped liability for negligent physicians who were independent contractors. That changed in *Darling* v *Charleston Cmty. Hosp.*[25], where the Illinois Supreme Court held a hospital liable for a surgeon's malpractice on grounds that hospital management of its facilities imposed a duty to patients. Many similar cases followed (often termed corporate negligence), expanding hospitals' liability for the acts of staff. Some courts then experimented with a more sweeping enterprise model. For instance, proposals of exclusive hospital enterprise liability would make hospitals answer for any patient injury occurring in their premises due to any provider, and even bar suit against individual caregivers.[26] Peters has described this idea, arguing that "the hospital enterprise liability would make hospitals liable for all patient injuries occurring in the hospital that are the product of provider negligence" and thus align incentives toward system-wide safety improvements.[27] This wave of expansive doctrine inspired advocates and critics alike: advocates saw it as a powerful deterrent and loss-spreading mechanism, while critics (and some courts) resisted.

By the late 20th century, many jurisdictions had curtailed or rejected broad enterprise liability. For example, courts often held that holding an insurer or parent corporation liable as part of an enterprise was inequitable, especially where legislature had allocated risk differently.[28] In the hospital context, narrow corporate-negligence standards prevailed: hospitals could be liable for negligent credentialing or supervision, but not as blanket insurers for all physicians' acts. Thus, the enterprise concept beyond those bounds had limited traction. In other areas, analogies to enterprise liability (e.g. joint and several liabilities among co-conspirators) remained rare or explicitly disavowed. Some jurisdictions incorporated enterprise theory into tests for piercing the corporate veil (treating affiliated companies as a single entity), as in *Mortimer* v *McCool*,[21] but that is a corporate-law construct distinct from tort doctrine.

Notably, most of this case law predates modern AI issues, and critics remain circumspect about applying enterprise liability automatically. Contemporary critics insist enterprise liability must satisfy both fairness and efficiency criteria to justify overriding individual fault principles, and many concluded it often does not.[15] In practice the doctrine has been limited by concerns over "swallowing up" all individual responsibility, line-drawing problems, and subverting long-standing duties (e.g. between physicians and patients).[29]

It is fair to say that, absent explicit statutory reform, pure enterprise liability has seen only sporadic success in court, often giving way to narrower fault-based or statutorily defined liability schemes.

The cluster of cases referred below will supply a compact jurisprudential toolkit for thinking about how tort law can and should respond to harms produced by algorithmic systems. Analysing these authorities together clarifies two related propositions. First, courts already possess a set of doctrines that can be adapted to close many of the accountability gaps created by AI. Second, the history of corporate personhood, as highlighted *in Santa Clara County* v *Southern Pacific Railroad Co.*[30] counsels us to treat personhood instrumentally i.e. the law creates legal subjects for pragmatic allocation of legal incidents, not out of metaphysical necessity. That dual lesson undercuts any simple narrative that the only way to regulate autonomous systems is to grant them independent legal personhood. The hospital cases, with Darling case at the centre, are important because they show how courts have historically translated systemic control into institutional duty. In Darling the court refused to permit hospitals to hide behind the doctrinal shield of independent contractor status when hospital organisation and practices materially shaped the safety of patients. The decision did not invent a new category; rather, it recognised that an institution that presents medical services to the public owes enforceable duties. This is conservative in the sense that it preserves traditional tort concepts, but it is also fits technological change. The lesson is practical: the law can reach organisational choices that create systemic risk without inventing new categories of legal capacity.

Sindell v *Abbott Laboratories*[29] supplies the complementary doctrine for cases in which causal indeterminacy frustrates an individualized fault inquiry. Sindell confronted a classic market-share problem: the plaintiff could not identify which manufacturer produced the harmful drug taken years earlier. The court's market-share allocation and enterprise reasoning were logical to avoid leaving victims remediless. Applied to AI, this logic legitimises a proportional or collective allocation

when opacity and distributed supply chains make attribution impossible. Sindell thus offers a doctrinal template for limited, proportionate collective liability in AI cases, provided careful limits are respected so that liability does not become arbitrary.

The history of corporate personhood often traced through *Santa Clara County* v *Southern Pacific Railroad Co.*[31] is salient because it frames the law's capacity to create legal subjects. Corporate personhood was adopted for instrumental reasons: to simplify transactions, to assign rights and duties to a collective that acts through human agents. This history matters for AI debates because it shows that the law need not endow a machine with full personhood to allocate legal incidents effectively. Kurki's "bundle" conception of personhood, which treats personhood as an assemblage of legal incidents allocated for functional ends, dovetails with this caution. The better course in most tort settings is to allocate liabilities, duties to warn, insurance obligations and enforcement tools to the human and corporate actors who design, deploy and profit from AI systems exactly the instrumental use of personhood the corporate example illustrates rather than to confer autonomous legal status on algorithmic code.

Platform cases and incidents bring these doctrinal lessons into contemporary relief. *Liu* v *Uber Technologies, Inc.*[32] and the Elaine Herzberg Uber incident[33] highlights the liability vacuum that arises when platforms disclaim operational control. Geisser's critique of Uber rests on the observation that driver's contractor status masks functional control of the algorithm. In Liu case, the driver logged in, system-dictated behaviours, and disputed timing of "on-trip" status show how technical categories can be used to avoid liability. The UK Supreme Court's decision in *Uber BV* v *Aslam*[34] confirms that courts may look beyond contractual labels to economic reality and control. Aslam vindicates the proposition that a platform's algorithmic governance can, in substance, create employment-like obligations and responsibilities. For AI governance, these holdings mean that courts are willing to treat algorithmic command structures as legally significant indicia of control that trigger institutional duties and vicarious responsibility where the facts support it.

Product and misrepresentation litigation also remains crucial. The *Young* v *Tesla, Inc.*[15] litigation cases reveal that where an AI system can be reasonably characterized as a "product" or where marketing misstates capabilities, established product-liability and consumer-protection doctrines can reach wrongdoing. These doctrines are important because they do not require proving state-of-the-art foreseeability in the same way negligence does; a defect or a misleading representation can establish liability even when causal chains are complex. But the automated vehicle litigation also demonstrates limitations like continuous learning and updates.

Taken together, these authorities point to a calibrated architecture rather than a single doctrinal leap. Where an institution controls and benefits from AI deployment as Darling and Aslam suggest courts should and do impose duties tailored to that control. Where causation is genuinely indeterminate and the industry collectively produced the risk, Sindell-style allocation provides a defensible, limited method of sharing loss. Where a marketed product or representation is at fault, product and consumer protection law remain powerful tool. At the same time, the corporate-personhood tradition and Kurki's bundle theory invite a pragmatic distribution of legal incidents. It gives liability wherever control and capacity to prevent harm rest, and attach insurance, disclosure and governance obligations to the human organisations that can internalise risk.

Granting full legal personhood to algorithms may create a legal anomaly that protects the human enterprises that actually design, profit from and can regulate them. The historical pattern shows that the law has preferred to translate moral expectations into duties and liability rules rather than to reify non-human agents as independent legal subjects. That approach follows tort law's corrective aims like compensation and deterrence. The path for policymakers, therefore, is doctrinal creativity within the human-centred architecture. Such as, extending institutional duties where control warrants, enterprise allocations and product-liability remedies for clear defects.

Enterprise Liability in the AI Context

Emerging scholarship suggests enterprise liability may have renewed relevance for harms from AI systems. AI technologies, especially opaque "black-box" systems, blur lines of causation and predictability. As Chan observes, AI systems can involve "algorithmic designers, software engineers, data curators, hospitals, and clinicians," each making distinct decisions, so that no single party clearly controls risk throughout the product cycle.[35] Existing frameworks often prove inadequate: if a physician relies on an AI recommendation that turns out harmful, is the doctor negligent for trusting it, or is the programmer negligent

for a coding flaw? Courts have noted that without a principled way to apportion responsibility among the various players, "patients may find it difficult to recover" and defection of one player can cause gaps.[18] Against this backdrop, some academics propose a common enterprise approach to AI liability. For example, Chan recommends treating all "participants in the health care machine-learning ecosystem" providers, manufacturers, and deploying institutions as a single enterprise. Under this joint-enterprise theory, if an AI system injures a patient, "all groups involved in the use and implementation of the AI system should jointly bear some responsibility," even if fault cannot be pinpointed to a specific actor.[36] This mirrors David Vladeck's suggestion of an enterprise liability rule: when an injury is AI-caused, instead of proving fault, one infers liability on the actors behind the system.[15] The advantage is that injured parties could be compensated without needing to crack the "black box" to find a negligent defect. In Chan's view, treating a hospital, its clinicians, and the AI provider as co-enterprises incentivizes them collectively to ensure safety and aligns with the need for system-level oversight.[18] In short, applying enterprise theory to AI would be a deliberate allocation of risk and costs among all beneficiaries of the AI technology, shifting from a narrow tort to a broader cost-spreading regime.[37]

Several authors similarly note that traditional doctrines struggle with AI's novelty. Sullivan and Scott observes that current models may be "insufficient to address the realities" of AI-driven injury.[38] It highlights enterprise liability as one possible solution alongside adjusting standards of care. This approach remains largely hypothetical, but it is gaining traction in academic discussion as a way to bridge the "responsibility gap" created by autonomous systems.

Critiques and Challenges

Enterprise liability has been criticized for stretching tort law. Civil recourse theorists in particular argue that tort's purpose is to vindicate rights and redress recognized wrongs, not to serve as a broad regulatory cost-spreading scheme.[39] Because traditional tort requires a breach of a duty owed to the injured party, relaxing that link undercuts the very notion of a "wrong." Under pure enterprise liability, a person could have no tort claim if no individual actor owed or breached a duty to him, yet somehow the enterprise as a whole would pay. Critics view this as incoherent with tort's corrective-justice foundations.[17]

Under practical challenges, determining what constitutes the "enterprise" is difficult. For clinical AI, does the enterprise include all users of a software package, each hospital's IT department, or only a manufacturer/provider consortium? Keating acknowledges that the boundary between individual acts and collective activities is "not always clear," raising line-drawing problems.[38]

Enterprise liability also raises allocation issues: once an enterprise is found liable, how are damages divided among its members? Traditional joint-and-several rules might apply, but that could be inequitable. Moreover, subjecting broad groups of actors (who may not have control over each other's conduct) to joint liability may deter socially valuable activities or give rise to litigious discovery battles to split blame. Policy critics also worry that enterprise liability undercuts incentives to avoid accidents. If all members share the loss regardless of fault, each member's individual incentive to be careful diminishes. While some loss spreading can improve risk distribution, over-extending it risks moral hazard. Indeed, enterprise liability proposals often rely on administrative or contractual controls between participants (joint insurance, shared monitoring) that may not materialize spontaneously. Finally, imposing liability on powerful institutions (entire industries or consortia) raises concerns of fairness: smaller participants might resent subsidizing risks they did not cause. In sum, while enterprise liability offers a systematic way to address the diffusion of responsibility posed by AI, it departs significantly from tort's core doctrines. It requires rethinking fundamental questions of duty and blame. As one scholar concludes, enterprise liability reform "fails to remain faithful to the foundational principles of tort law" and may be best addressed by alternative solutions like no-fault schemes rather than stretching tort beyond recognition.[40]

Risk Society and AI Liability

The European Union's proposed Artificial Intelligence Act (AIA) represents a major shift in how law engages with emerging technology. Instead of waiting for harm to occur and then assigning blame through tort law, the AIA seeks to prevent harm through structured, anticipatory governance. In this context, the AIA can be read not only as a piece of regulatory legislation but also as a legal response to the broader social challenge of managing technologically produced risks that no longer have a single, identifiable cause.[41]

Beck's risk society explains that as societies advance technologically, they generate new forms of "manufactured uncertainty." These are risks that emerge not from natural causes but from human innovations like industrial systems, biotechnology, and now artificial intelligence.[42] The design, training, and deployment of AI involve multiple layers of human and algorithmic agency, making it difficult to locate individual fault when harm occurs. Traditional tort law, based on negligence and direct causation, struggles with such complexity. The EU's risk-based model acknowledges this limitation. It shifts the emphasis from who caused the harm to how the harm can be prevented, adopting a preventive, system-oriented approach to liability.[24]

The AIA classifies AI systems according to their level of risk: unacceptable, high, limited, and minimal. The law also mandates corresponding safety and oversight mechanisms. This ex ante model of risk regulation contrasts with the ex post logic of tort law, which operates after damage occurs. It introduces detailed compliance duties such as data governance, human oversight, and documentation. However, as several scholars have noted, these duties risk becoming procedural rather than substantive. If responsibility is defined only by following regulatory steps, real accountability may disappear. Beck calls this phenomenon "organised irresponsibility": a condition where risks are managed through institutional compliance, yet no one is morally or legally answerable when harm occurs.[43]

From a tort law perspective, this proceduralisation of responsibility creates a significant gap. Regulation prevents harm, but tort law delivers justice when prevention fails. If compliance with the AIA provides a complete defence against liability, victims may find no route to compensation. Therefore, regulatory and tort frameworks should not be treated as substitutes but as complementary systems. Compliance should provide presumptive evidence of care, not complete immunity. This preserves tort law's corrective and distributive functions, ensuring that victims are not left unprotected within the bureaucratic machinery of risk governance. One way to integrate the preventive and compensatory dimensions of law is through enterprise liability. This doctrine recognises that modern harms often arise from collective rather than individual activity. In Sindell case for example, the court held several drug manufacturers collectively liable when the specific source of the injury could not be proven. The same principle can be applied to AI systems, where multiple entities like developers, data providers jointly create risk. Enterprise liability ensures that those who profit from AI also bear the costs of its failures. It resonates with Beck's vision of collective risk internalisation, where responsibility is shared in proportion to control and benefit.

Institutional duty cases such as Darling and Uber BV reinforce this logic. In Darling, the hospital was held responsible for systemic failures in supervision, even though the negligent actor was an independent doctor. Similarly, Aslam established that Uber exercised significant control over drivers through its algorithmic management system, making it functionally an employer. These decisions show that law can impose duties where institutions have structural control over risk, regardless of formal contractual distinctions. Applying this reasoning to AI means that corporations cannot deflect blame onto autonomous systems. If they design, deploy, or profit from AI, they should be legally bound to monitor and mitigate its harms.

The AIA's transparency and documentation requirements could serve as crucial links between regulation and tort law. If properly implemented, they would not only support compliance but also provide evidence for liability proceedings. Records of training data, model performance, and decision logic can help courts trace causation and apportion liability. However, these mechanisms must function as tools of accountability, not as symbolic procedures. Risk management must remain dynamic: AI systems evolve through updates and retraining, so monitoring and audit duties should continue throughout a system's life cycle.[44] To further strengthen accountability, additional measures may be necessary. Mandatory insurance or collective compensation funds could ensure that victims receive redress even when individual responsibility cannot be precisely determined. Such mechanisms mirror existing frameworks in environmental and medical liability, where risks are diffuse and cumulative. They also reflect Beck's insight that in a risk society, residual risks must be institutionally managed rather than left to chance. The deeper significance of the AIA lies in its attempt to translate the conditions of the risk society into law. It accepts that technological risk cannot be eliminated but can be managed. Yet this management must not come at the expense of moral responsibility. Tort law retains a vital role as the system that personalises accountability within a world of structural

risks. Enterprise liability and institutional duty doctrines can help ensure that accountability remains meaningful and human-centred.

In conclusion, the EU's risk-based approach to AI regulation represents both progress and challenge. It reflects an evolved understanding of risk, one that recognises the systemic nature of technological harm, but it also risks transforming responsibility into procedural compliance. To prevent this, tort law and regulation must operate in tandem: regulation to anticipate and reduce risk, and tort law to correct and compensate when risk becomes harm. A balanced framework anchored in collective responsibility, institutional duty, and transparent oversight can bridge the gap between preventive governance and corrective justice. In doing so, it allows the law to manage AI risk without losing sight of the human values that justify legal responsibility in the first place.

Conclusion

The evolution of artificial intelligence has exposed fundamental weaknesses in traditional liability frameworks by creating modes of action, decision-making, and risk that no longer align with the core assumptions of tort law. Across the scholarship examined, a consistent theme emerges: AI distributes agency across developers, deployers, datasets, and algorithms in ways that render single-actor fault attribution inadequate. Incidents such as the Uber autonomous vehicle crash illustrate this gap vividly. The result is a liability vacuum in which victims lack clear routes to redress, and accountability is diluted among many actors who can plausibly claim compliance.

Ulrich Beck's risk society theory helps explain this structural shift. AI creates "manufactured risks" like systemic, cumulative, and deeply embedded in socio-technical architecture. As the reviewed articles emphasise, these risks challenge the adequacy of classical tort doctrines such as negligence, vicarious liability, or even strict product liability. The opacity and learning capacity of AI systems weaken traditional tools of causation, foreseeability, and defect analysis. At the same time, the procedural structure of the EU's risk-based Artificial Intelligence Act shows that regulation alone cannot restore accountability. While the Act's classification, documentation, and transparency duties provide important preventive safeguards, they risk transforming responsibility into bureaucratic compliance. Scholars warn that if regulatory conformity becomes a shield against liability, the system will reproduce "organised irresponsibility," where harm is absorbed socially while developers and deployers escape meaningful scrutiny.

The literature points toward the need for an integrated solution rather than a replacement of tort by regulation or vice versa. Building on these insights, this paper proposes a hybrid liability model that draws from both preventive and corrective approaches. The preventive component mirrors the EU's risk-based regulation: mandatory risk assessments, data-quality controls, human oversight requirements, and robust documentation. These obligations ensure that organisations remain aware of the risks they create and actively work to mitigate them before deployment. However, regulatory compliance under this model does not extinguish liability, it merely informs it.

The corrective component preserves tort law's essential functions. Developers, deployers, and institutional users remain liable where they fail to meet standards of reasonable care or where they profit from activities that generate systemic risks. Enterprise liability plays a central role here. When harm emerges from interconnected actors, liability should be allocated across those who exercised control or derived economic benefit from the system. This prevents the burden of technological risk from falling solely on victims and discourages opportunistic fragmentation of responsibility within the AI supply chain.

Institutional duty doctrines further support this structure by recognising that organisations bear responsibility when they shape or manage decision-making environments, whether through human supervisors or algorithmic governance tools. AI deployers therefore carry ongoing duties of oversight, not only at the time of design or procurement but throughout the system's operational life. Because AI systems evolve, these duties include continuous monitoring, regular audits, transparent reporting, and periodic reassessments of risk.

To address residual, unpreventable harms, particularly those arising from complex, adaptive systems, the hybrid model incorporates collective risk-management instruments. Mandatory insurance, compensation funds for high-risk sectors, and pooled industry mechanisms ensure that victims are compensated even when precise fault cannot be established. These measures reflect the broader academic consensus that modern technological risks must be internalised within the industries that profit from them, rather than externalised to individuals and the public.

Taken together, the doctrinal insights from tort law, the sociological analysis of risk society, and the regulatory innovations of the EU converge on a shared conclusion: AI governance requires a framework that views responsibility as both preventive and corrective, both individual and collective. The hybrid model advanced in this paper captures this duality. It offers a pathway that preserves accountability, encourages innovation, and builds a coherent legal response to the distinctive challenges posed by artificial intelligence. By integrating risk-based regulation with an expanded set of liability principles, the law can respond not only to the technical complexity of AI but also to the deeper social transformations it represents.

References

1 https://www.bbc.com/news/technology-54175359 (accessed on11 September 2025).

2 Ahuja V K, Artificial intelligence and copyright: Issues and challenges, *ILI Law Review*, WinterIssue (2020) 270.

3 Zohuri B & Behgounia F, *Application of artificial intelligence driving nano-based drug delivery system*, In Philip A, Shahiwala A, Rashid M, Faiyazuddin M, *A handbook of artificial intelligence in drug delivery*, (Academic press, Cambridge), 2023, 145-49.

4 Feinberg J, Duties, rights, and claims, *American Philosophical Quarterly*, 3 (2) (1966) 137.

5 Zimmerman, E J, Machine minds: Frontiers in legal personhood, *SSRN Electronic Journal*, (2015) 2.

6 https://www.diva-portal.org/smash/get/diva2:1115160/FULLTEXT01.pdf (accessed on 22 January 2025).

7 https://www.niti.gov.in/sites/default/files/2023-03/National-Strategy-for-Artificial-Intelligence.pdf (accessed on 25 January 2025).

8 https://www.lexology.com/library/detail.aspx?g=6ee25b10-23df-4452-b138-4da1104f70a7 (accessed on 24 January 2025).

9 Kelsen H, *General theory of law and State*, (Routledge, London), 1 (2017) 157.

10 https://www.sci.brooklyn.cuny.edu/~schopra/agentlawsub.pdf (19 December 2024).

11 Fauzan M P N, Amarta D, Tobias E, Ricardo V & Melania G F, Wandering with artificial intelligence and its obscure legal liability, *Indonesia Law Review*, 11 (2) (2021) 169.

12 Montagnani L, M & Cavallo M, Liability and emerging digital technologies: An EU perspective, *Notre Dame Journal of International & Comparative Law,* 11 (2) (2021) 208.

13 Geisser L, Risk, reward, and responsibility: A call to hold Uberx, Lyft, and other transportation network companies vicariously liable for the acts of their drivers, *Southern California Law Review*, 89 (2015) 317.

14 Kurki V, *A Theory of Legal Personhood* (Oxford University Press, Oxford) 2019, 241.

15 Chan B, Applying a common enterprise theory of liability to clinical AI systems, *American Journal of Law & Medicine,* 47 (4) (2021) 351.

16 H -Y Chiu I & Lim E, Managing corporations' risk in adopting artificial intelligence: A corporate responsibility paradigm, 20 *Washington University Global Studies Law Review,* 20 (2021) 349.

17 Keating C G, The theory of enterprise liability and common law strict liability, *Vaderbilt Law Review*, 54 (3) (2001) 1292.

18 Sullivan R H & Schweikart J S, Are current tort liability doctrines adequate for addressing injury caused by AI, *AMA Journal of Ethics,* 21 (2) (2019) 160.

19 Witt F J, Speedy Fred Taylor and the ironies of enterprise liability, *Columbia Law Review,* 103 (1) (2003) 1.

20 See B, Paging Doctor Robot: Medical artificial intelligence, tort liability, and why personhood may be the answer, *Brooklyn Law Review,* 87 (1) (2021).

21 Keating C G, Fairness and two fundamental questions in the tort law of accidents (working paper), *SSRN Electronic Journal*, (2000) 1.

22 Keating C G, The idea of fairness in the law of enterprise liability, *Michigan Law Review*, 95 (5) (1997) 1266.

23 Hallevy G, Unmanned vehicles: Subordination to criminal law under the modern concept of criminal liability, *Journal of Law, Information and Science,* 21 (200) (2012) 1.

24 Maher D, Keeping pace with advancing artificial intelligence in healthcare: A call for increased regulation and legislative action to avoid the deflection of liability, *Annals of Health Law Advance Directive,*31 (2021) 165.

25 Peters G P, Resuscitating hospital enterprise liability, *Missouri Law Review*, 73 (2008) 369.

26 Hallevy G, The criminal liability of artificial intelligence entities - From science fiction to legal social control, *Akron Intellectual Property Journal*, 4 (2) (2010) 171.

27 *Gopal* v *Kaiser Found. Health Plan, Inc*. B259808 (Cal. Ct. App. 2d Dist. May 26, 2016)

28 255 A.3d 261 (Pa. 2021),

29 118 U.S. 394 (1886).

30 26 Cal. 3d 588.

31 551 F.Supp.3d 988 (N.D. Cal. 2021).

32 *State of Arizona* v *Rafael(a)Vasquez*, No. CR2020-001853-001 (Ariz. Super. Ct. July 5, 2021).

33 [2021] UKSC 5.

34 1:21-cv-00917, (D.N.M.).

35 Hodge D S Jr, The medical and legal implications of artificial intelligence in health care - An area of unsettled law, *Richmond Journal of Law & Technology,* 28 (3) (2024) 405.

36 Tonti N, Who to blame? AI or your physician, *Quinnipiac Health Law Journal*, 27 (2) (2024) 261.

37 Jessica S A, From jeopardy! to jaundice: The medical liability implications of Dr. Watson and other artificial intelligence systems, *Louisiana Law Review,* 73 (4) (2013) 1049.

38 Tilley C C, Tort Law inside out, *Yale Law Journal,* 126 (2017) 1320.

39 Sword M, To err is both human and non-human, *University of Missouri-Kansas City Law* Review, 88 (2019) 211.

40 Baş F, Artificial intelligence, human and society in the context of Ulrich Beck's risk society theory, *Journal of Eskişehir Osmangazi University Faculty of Theology*, 12 (2025) 43.

41 Beck U, World risk society as cosmopolitan society? Ecological questions in a framework of manufactured uncertainties: ecological questions in a framework of manufactured uncertainties, *Theory, Culture & Society,*13 (4) (1996) 1.

42 Hallevy G I & Robot-I, Criminal: When science fiction becomes reality: Legal liability of AI robots committing criminal offenses, *Syracuse Science & Technology Law Reporter* 1 (2010) 1.

43 Frank X, Is Watson for oncology per se unreasonably dangerous? Making a case for how to prove products liability based on a flawed artificial intelligence design,*American Journal of Law & Medicine*, 45 (2019) 273.

44 Chamberlain J, The risk-based approach of the European Union's proposed Artificial Intelligence Regulation: Some comments from a Tort Law perspective, *European Journal of Risk Regulation*, 14 (2023) 1.