International Journal of Multidisciplinary Research and Growth Evaluation.

Legal Liability for Harm Caused by Artificial Intelligence Systems

Pham Van Vinh
Ho Chi Minh City Bar Association, Ho Chi Minh City, Vietnam

* Corresponding Author: **Pham Van Vinh**

Article Info

ISSN (Online): 2582-7138
Volume: 07
Issue: 05
Received: 21-07-2026
Accepted: 19-08-2026
Published: 17-09-2026
Page No: 577-581

Abstract
This article examines the attribution of liability when harm arises in the development, provision, and deployment of artificial intelligence (AI) systems, with particular emphasis on civil liability for damages. Through doctrinal analysis, engagement with scholarly arguments, and a comparison between Vietnamese and European Union (EU) law, it argues that the autonomy of AI does not, by itself, create a new legal person to bear liability. The central difficulties instead lie in the fragmentation of control over risk and the evidentiary imbalance between injured parties and actors in the AI value chain. Vietnam's Law on Artificial Intelligence 2025 establishes a compensation obligation for deployers in certain cases involving high-risk AI systems. Nevertheless, it does not fully resolve the problems of proving causation or allocating the ultimate financial burden when the source of harm lies in development or supply. This article proposes a layered framework: an identifiable point of recourse for victims of certain high-risk activities; the continued application of liability grounded in conduct, defects, and control over risk; and statutory rules on evidence preservation, access to evidence, and recourse against responsible parties. These proposals concern legal reform rather than rights and obligations already comprehensively established under current law.

DOI: https://doi.org/10.54660/.IJMRGE.2026.7.5.577-581

Keywords: artificial intelligence, compensation for harm, civil liability, product liability, causation, risk allocation

1. Introduction

Consider a hospital that deploys an AI system to assist in interpreting diagnostic images. The system fails to detect an indication of disease; a physician relies on its output; and treatment is delayed. Determining liability requires more than asking what output the AI generated. It also requires examining whether the input data were appropriate, whether the provider disclosed known limitations, whether the hospital validated the system under actual operating conditions, and how carefully the physician assessed its output. This hypothetical illustrates how a single injury may be connected to decisions made by different people at different stages. Research on AI liability likewise shows that machine learning, interconnected systems, and multi-actor supply chains can shift control over risk away from end users and towards developers, providers, or data suppliers (Zech, 2021) [11].

The research question is therefore not whether AI can be "at fault" in the same sense as a person. Rather, it is which actor should bear the adverse legal consequences, on what legal basis, and subject to what standard of proof. In Vietnam, this question must be considered in light of Law No. 134/2025/QH15 on Artificial Intelligence, which entered into force on 1 March 2026. The statute distinguishes developers, providers, deployers, and users, while Article 29 directly addresses compensation. It would consequently be inaccurate to suggest that Vietnam has no rules at all governing liability in connection with AI (National Assembly, 2025) [9].

This article argues that traditional principles of civil liability remain indispensable, but they cannot be applied simply by treating every AI output as the wrongful act of a particular person. Protecting injured parties without releasing those who control other stages of risk requires separating two questions: who should compensate the victim in the first instance, and who should

ultimately bear the cost? That distinction provides the basis for assessing both Article 29's limitations and the relevance of EU experience.

2. Theoretical Foundations of Liability for AI-Related Harm

Under Article 3(2) of Vietnam's Law on Artificial Intelligence 2025, an AI system is a machine-based system that operates with varying levels of autonomy, may exhibit adaptiveness after deployment, and infers from inputs how to generate predictions, content, recommendations, or decisions that can influence physical or virtual environments. For liability purposes, the decisive issue is not the label "AI," but how the system participates in the chain leading to harm. It may directly control a device, recommend a course of action to a human decision-maker, or generate information that an organisation passes on to others (National Assembly, 2025) [9].

The phrase "harm caused by AI" in the title is shorthand. Legally, at least three situations must be distinguished. First, a system's output or operation may directly affect a person or property. Second, a person may act on an AI output, and that action may result in injury. Third, harm may arise in a process that uses AI even though an independent factor is the actual cause. The presence of AI in a process does not, by itself, prove causation. Conversely, an intervening human decision does not necessarily sever the connection between the system's design and the resulting harm.

Legal liability may be administrative, criminal, or civil. This article focuses on the obligation to compensate for harm. Within that inquiry, *causal responsibility*, meaning that a factor contributed to an outcome, must be distinguished from *liability to compensate*, meaning that the law assigns the resulting burden to an actor under specified conditions. Hage (2017) [5] argues that autonomous agents could, in theory, be treated as responsible for their own actions; the crucial question is whether doing so would be desirable. This argument usefully challenges the assumption that legal responsibility can only ever be attributed to human beings. Yet theoretical possibility does not establish the practical value of making AI a bearer of liability. If a system has no independent assets and assigning responsibility to it weakens incentives for the organisations behind it to prevent harm, victims may be worse off. Vietnamese law presently assigns rights and obligations to organisations and individuals engaged in AI activities rather than establishing AI systems as independent persons liable for compensation (Hage, 2017; National Assembly, 2025) [5, 9].

Autonomy and post-deployment learning affect liability in a more specific way: a developer may be unable to anticipate every output, while a deployer may lack control over the model's internal structure. The "black box" problem is primarily one of obtaining and interpreting evidence. An injured party may prove that an adverse decision was made but remain unable to tell whether it arose from the data, the model, its integration into a wider system, or its use. Padovan, Martins, and Reed (2023) [10] argue that explainable AI can help courts apply existing liability frameworks. That insight is valuable, but a technical explanation of an output must not be equated with a legal finding about the cause of injury. The latter also requires an assessment of each actor's duties and alternative explanations (Padovan *et al*., 2023; Zech, 2021) [10, 11].

The appropriate theoretical starting point is therefore not a rule that the person closest to an AI system must bear every resulting loss. A defensible framework should perform three functions at once: offer injured parties a practicable route to compensation, place the costs of risk on actors able to prevent it, and preserve sufficiently rigorous standards of proof so that every adverse event occurring alongside AI use is not automatically attributed to AI.

3. Challenges in Identifying and Allocating Liability

Conduct-based and fault-based approaches are persuasive where a specific deficient decision can be identified: a developer omits necessary testing, a provider fails to disclose a known limitation, a deployer uses the system beyond its intended scope, or a user disregards a clear warning. Such approaches encourage each actor to discharge duties within its sphere of control. However, the claim that "no one could predict the exact output, so no one was at fault" misidentifies what must be foreseeable. A duty of care need not require prediction of each individual result; it may instead require anticipation of a *category of risk*, appropriate testing, warnings, monitoring, and intervention when anomalies emerge. The harder questions concern the applicable standard of care at the relevant time and access to the evidence needed to assess compliance (Zech, 2021; Chen, 2026) [11].

At the same time, every malfunction should not be attributed to the deployer merely because it introduced the system into use. Chen (2026) [1] advocates a deployer-centred approach because the deployer determines the operational setting. The same study acknowledges a limitation, however: a deployer accessing a service through a closed application programming interface may have little visibility into the provider's data, parameters, or safety mechanisms. Zech (2021) [11], by contrast, emphasises that liability should track the ability to control risk, including when that ability lies with a manufacturer or infrastructure operator. Read together, these views show why a fixed role label cannot resolve the whole question. Proximity to the victim may make an actor a suitable initial point of recourse, whereas substantive control over risk is an important basis for allocating the ultimate burden (Chen, 2026; Zech, 2021) [1, 11].

Causation becomes particularly difficult in the hypothetical diagnostic case. If an image is flawed, the model responds inaccurately, and the physician fails to notice the problem, the final outcome cannot simply be attributed in full to one stage. The sequence must be reconstructed, including the model version and any updates, the actual input, warnings displayed, opportunities for human intervention, and the counterfactual question whether compliance with a particular duty would have avoided or reduced the harm. This is a method of analysing individual cases, not a claim that existing law automatically treats every event in the sequence as a legal cause. Padovan and colleagues (2023) [10] emphasise explainability as an aid to attribution, while Zech (2021) notes that even liability without fault does not dispense with causation (Padovan *et al*., 2023; Zech, 2021) [10, 11].

Product liability addresses some difficulties associated with fault. An injured party need not necessarily establish that the producer was negligent, but must establish a defect, harm within the protected scope, and a causal connection. With AI, the meaning of "defect" is itself contestable. A disclosed error rate does not necessarily establish safety in every

setting, while a single incorrect output does not, without more, prove a design defect. The assessment must take account of intended use, reasonably foreseeable use, warnings, the level of safety that may legitimately be expected, and changes after deployment. The EU's revised Product Liability Directive specifically addresses products capable of learning or acquiring new features after being placed on the market. It also provides for access to evidence and evidentiary presumptions under defined conditions (European Parliament & Council of the European Union, 2024b) [4].

A regime of liability irrespective of fault for high-risk activities has a different advantage: a deployer cannot shift the entire cost of an accident to the victim merely by showing compliance with applicable procedures. Yet *liability irrespective of fault is not unconditional liability*. It remains necessary to establish that the harm falls within the regime's scope and to consider statutory defences. Moreover, placing all costs on a deployer when the underlying problem arose in a model controlled by someone else may distort incentives to prevent harm. Zech (2021) [11] therefore argues for attention to both producers and operators and discusses insurance for risks that are difficult to attribute to a single actor. This is a policy consideration, not proof that a common insurance fund would work in every sector (Zech, 2021) [11].

These limitations support a two-level approach to allocation. At the external level, the law may identify a relatively accessible party that must compensate victims of certain high-risk AI activities. At the internal level, a court should assess the contribution and control of developers, providers, integrators, deployers, users, and parties who intervene unlawfully. Joint and several liability should arise only where the statutory requirements for jointly caused harm or another specific basis of joint liability are met. It cannot be inferred merely because several actors participated in the AI value chain. This structure facilitates compensation without turning the deployer into the insurer of every risk created by others.

4. EU Experience and Its Relevance to Vietnam

The EU illustrates the need to distinguish *ex ante risk regulation* from *ex post compensation*. Regulation (EU) 2024/1689, commonly called the AI Act, imposes duties on actors in the AI value chain. For high-risk systems, these include requirements concerning risk management, technical documentation, logging, and human oversight. Such duties may inform the standard of conduct and generate evidence when disputes arise, but the AI Act is not a general compensation code that automatically identifies who must pay whenever harm occurs. Its general date of application is 2 August 2026, although certain provisions applied earlier under the Regulation itself (European Parliament & Council of the European Union, 2024a) [3].

In product liability, Directive (EU) 2024/2853 brings software, including AI systems, within its definition of a product, subject to the Directive's conditions. It provides for liability irrespective of fault on the part of relevant economic operators where a defective product causes a protected type of harm. Article 9 allows a claimant with a plausible compensation claim to request disclosure of relevant evidence, while Article 10 provides for presumptions of defectiveness or causation in specified circumstances; these presumptions remain rebuttable. The Directive does not cover every AI-related loss. Pure economic loss, privacy violations, or discrimination do not, by themselves, trigger liability under this Directive, although other legal avenues may be available. The Directive has entered into force, but Member States have until 9 December 2026 to transpose it. Accordingly, its new framework should not be portrayed, at the time of writing, as already operating uniformly in national practice (European Parliament & Council of the European Union, 2024b) [4].

Another distinction is essential. The European Commission's 2022 proposal for an *AI Liability Directive*, which addressed evidence and fault-based liability, was withdrawn. The withdrawal appears in a notice published in the EU's Official Journal on 6 October 2025. The EU therefore does not offer a separate and comprehensive civil liability regime for all AI-related harm. Its current framework combines the AI Act, defective-product rules, and other applicable liability regimes (European Commission, 2025; European Parliament & Council of the European Union, 2024a, 2024b) [2, 3, 4].

The principal lesson for Vietnam lies in connecting three elements: duties to manage risks before harm occurs, access to evidence after it occurs, and rules allocating liability to those who could control the relevant source of risk. Vietnam should not transplant the EU's concept of a "product" or its evidentiary presumptions without considering differences in the scope of protection offered by the two legal systems. Vietnamese law already includes a specific rule concerning deployers of high-risk AI systems. The task is to make that rule work alongside civil, procedural, and consumer protection law rather than to create an overlapping liability regime with unclear boundaries (National Assembly, 2025; European Parliament & Council of the European Union, 2024b) [9, 4].

5. Vietnamese Law and Directions for Reform

Current law. The Civil Code 2015 provides the foundation for non-contractual compensation. Article 584 sets out the grounds for liability, Article 585 establishes principles of compensation, and Article 587 addresses harm jointly caused by multiple persons. Resolution No. 02/2022/NQ-HDTP explains that cases under Article 584(1) require consideration of an infringing act, harm, and causation. It also addresses a claimant's ability to ask the court to collect evidence that the claimant cannot obtain independently. Article 601 of the Civil Code provides a special regime for highly dangerous sources. An AI system classified as "high-risk" under the Law on Artificial Intelligence does not, merely for that reason, automatically become a highly dangerous source under Article 601. The former is a regulatory classification; application of the latter still depends on that article's own subject matter and conditions (National Assembly, 2015; Judicial Council of the Supreme People's Court, 2022) [7, 6].

The Law on Protection of Consumer Rights 2023 also provides a basis for claims involving defective products or goods. Under Article 34, business organisations and individuals within its scope must compensate consumers for harm caused by defective products or goods they supplied, even if they did not know of or were not at fault for the defect, subject to the exemptions in Article 35. This rule cannot automatically be extended to *every* stand-alone AI service. A claimant must first establish whether the subject of the dispute is a product, goods, or a service, and which provisions apply. The statute also recognises compensation claims connected with services, but this does not mean every service claim falls under Article 34's defective-product regime (National Assembly, 2023) [8].

Most directly, Article 29 of the Law on Artificial Intelligence 2025 provides that organisations and individuals who violate AI-related law and cause harm must compensate for it in accordance with civil law. Where a high-risk AI system is managed, operated, and used in compliance with applicable rules but nevertheless causes harm, the deployer must compensate the injured party. After paying compensation, the deployer may seek reimbursement from the provider, developer, or another relevant party *if the parties have agreed to such an arrangement*. Article 29(3) provides exemptions from the liability established in paragraph 2. Paragraph 4 deals with unlawful intrusion, seizure of control, or interference by a third party, including joint liability of a deployer or provider whose fault allowed the interference to occur. Article 14 additionally establishes relevant duties concerning risk management, documentation, logging, oversight, and cooperation in incident response. Civil liability insurance is encouraged, rather than made a general statutory requirement under that provision (National Assembly, 2025) [9].

Gaps and interpretive questions. Article 29(2) has an important advantage: it prevents a deployer from defeating a claim solely by arguing that it complied with the rules for managing and operating a high-risk system. Nevertheless, three problems remain. First, a claimant must still establish a connection between the harm and the system's operation; the provision should not be read as dispensing with causation merely because harm occurred. Second, the boundary between *compliance despite harm* under paragraph 2 and *breach of a duty* under paragraph 1 requires clarification when several actors are involved in the same incident. Third, the express reference to reimbursement "if the parties have agreed" makes reallocation under paragraph 2 substantially dependent on contract. Without such an agreement, an automatic statutory right of recourse cannot be inferred from that paragraph. Equally, it would be too broad to conclude that developers and providers are exempt from every independent basis of civil liability (National Assembly, 2025; Judicial Council of the Supreme People's Court, 2022) [9, 6].

Evidence presents a further problem. Article 14 requires relevant documentation and logs to be retained for high-risk AI systems where necessary, but the existence of records does not, by itself, ensure that claimants can obtain them promptly, within an appropriate scope, and in a form that can be tested. Any procedure must balance business confidentiality, personal data protection, and the right to contest evidence. Articles 9 and 10 of Directive (EU) 2024/2853 illustrate a possible legislative approach: an initial claim must have a reasonable basis; orders to disclose evidence should be necessary and proportionate; and adverse presumptions should arise only under legally specified conditions, not automatically whenever a system is difficult to explain (National Assembly, 2025; European Parliament & Council of the European Union, 2024b) [9, 4].

The author's position and recommendations. In the near term, guidance on Article 29 should follow a consistent sequence of legal analysis. Decision-makers should identify the system and its risk classification, establish each actor's actual function rather than rely exclusively on contractual labels, determine the harm and the events leading to it, and only then select the applicable basis of liability, exemptions, and reimbursement rules. A case record should permit the reconstruction, insofar as necessary for the dispute, of the system version, post-deployment changes, inputs, warnings, intervention logs, and security incidents. This is a proposed approach to adjudication and guidance, not a claim that the Law on Artificial Intelligence already prescribes a comprehensive AI-specific procedural code (National Assembly, 2025; Padovan *et al*., 2023) [9, 10].

In the longer term, the legislature should consider adding a *statutory basis for recourse* alongside contractual arrangements. It would apply where a deployer has compensated a victim under Article 29(2) and subsequently proves that the harm resulted from another actor's breach of a safety or disclosure duty, or from a defect within that actor's sphere of control. Such a rule would need to specify evidentiary conditions, the method for assessing each actor's contribution, and its relationship with Article 587 of the Civil Code. It should not automatically make every participant in the AI value chain jointly and severally liable for every injury. The initial compensation mechanism could thereby protect victims while leaving room for the ultimate cost to move towards the party positioned to prevent the relevant risk (National Assembly, 2015, 2025; Zech, 2021) [7, 11].

As to evidence, reform could allow courts to order the production of relevant records, logs, and technical information held by a party, subject to suitable confidentiality safeguards. A rebuttable presumption could also be considered where a party possessing relevant evidence breaches a duty to preserve it or fails to comply with a lawful disclosure order. The presumption must be limited: claimants should still provide initial facts indicating harm and the system's material involvement, so that missing logs do not become the sole basis for every claim. Liability insurance should be assessed sector by sector, particularly where potential losses are substantial and risks can be priced with reasonable reliability, rather than imposed indiscriminately on all developers and AI applications. These are proposals derived from existing record-keeping duties, current liability rules, and EU experience; they are not generally applicable rules already in force in Vietnam (National Assembly, 2025; European Parliament & Council of the European Union, 2024b; Zech, 2021) [9, 4, 11].

6. Conclusion

When an AI system is implicated in harm, the liable party is not "the AI" merely because its output was closest to the resulting injury. Liability must be determined by the applicable legal grounds and each actor's control over risk. A deployer may be the initial party required to compensate under Article 29(2) of the Law on Artificial Intelligence in cases covered by that provision. Developers, providers, users, or third parties may also bear liability on grounds involving their conduct, a defect, unlawful interference, or other relevant legal bases. Traditional civil liability principles therefore remain applicable, but they are insufficient without clearer rules on evidence and the allocation of the ultimate burden. This article favours refining the existing framework according to risk and substantive control, rather than granting legal personality to AI or imposing a single liability regime on every AI system (National Assembly, 2025; Zech, 2021) [11].

References

1. Chen Z. Operational responsibility in AI governance: a user-centric liability framework. AI Ethics. 2026;6:306. doi:10.1007/s43681-026-01163-7.
2. European Commission. Withdrawal of Commission

proposals (C/2025/5423). Official Journal of the European Union. 2025.
3. European Parliament, Council of the European Union. Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence. Official Journal of the European Union. 2024.
4. European Parliament, Council of the European Union. Directive (EU) 2024/2853 of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC. Official Journal of the European Union. 2024.
5. Hage J. Theoretical foundations for the responsibility of autonomous agents. Artif Intell Law. 2017;25:255-271. doi:10.1007/s10506-017-9208-7.
6. Judicial Council of the Supreme People's Court of Viet Nam. Resolution No. 02/2022/NQ-HDTP of 6 September 2022 guiding the application of certain provisions of the Civil Code on non-contractual liability for damages. 2022.
7. National Assembly of Viet Nam. Civil Code, Law No. 91/2015/QH13 of 24 November 2015. 2015.
8. National Assembly of Viet Nam. Law on Protection of Consumer Rights, Law No. 19/2023/QH15 of 20 June 2023. 2023.
9. National Assembly of Viet Nam. Law on Artificial Intelligence, Law No. 134/2025/QH15 of 10 December 2025. 2025.
10. Padovan PH, Martins CM, Reed C. Black is the new orange: how to determine AI liability. Artif Intell Law. 2023;31:133-167. doi:10.1007/s10506-022-09308-9.
11. Zech H. Liability for AI: public policy considerations. ERA Forum. 2021;22:147-158. doi:10.1007/s12027-020-00648-0.

How to Cite This
Pham VV. Legal liability for harm caused by artificial intelligence systems. Int J Multidiscip Res Growth Eval. 2026;7(5):577-581. doi:10.54660/.IJMRGE.2026.7.5.577-581.

Creative Commons (CC) License
This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution NonCommercialShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work noncommercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.