Minnesota Journal of Business Law and Entrepreneurship

Volume 2026, No. 2

ISSN: 1540-3270

Synthetic Harm and Legal Accountability: Structuring Tort Law Responses to Generative AI Systems

Bhumireddy Sai Srinivas Reddy

BBA, LLB (Hons.), LLM (USA)

Indiana (USA) Bar Admission Pending

Abstract

Generative AI has come a long way from its humble beginnings as something that was merely bandied about in casual conversation among friends and colleagues. Today, it is ubiquitous and has become an integral part of every digital platform that we use and interact with. The extent to which it has become an integral part of our lives is raises significant normative and legal concerns. While its proliferation is certainly having many positive effects on our lives, it is evident that there are many problems that come along with it, particularly in the context of tort law principles in the United States. While it is evident that there is an element of uncertainty that is associated with Generative AI and its outputs, can we really rethink and reapply some of the very old principles of tort law to this new and complex world? This paper is based on three major legal disruptions that have been caused by generative AI. First of all, the unpredictability of generative AI makes it hard to apply principles of negligence law because it is based on foreseeability and a quantifiable standard of care. Secondly, tort law is based on the concept of one person being liable for an injury. On the other hand, generative AI is based on complex systems that spread liability to developers, companies that use the technology, and the technology itself. Thirdly, there is difficulty in applying product liability law to generative AI because it is dynamic and is changing all the time through interactions with users and is thus considered to be either a product or a service. Furthermore, there is an overwhelming amount of information on platforms such as Widely used generative AI platforms and that is seeping into regular social media which is making it increasingly difficult to differentiate between what is real and what is artificial. This is posing a threat to legal principles that depend on clearly identifiable communicators and intent. Instead of seeking to fundamentally transform the law of torts, this paper proposes a more pragmatic approach to developing a parallel system that can be applied in conjunction with existing rules. This might include risk-based categorization, ex-ante duties, and ex-post accountability. The intention is to address the current issues while also being flexible to address future changes in generative AI.

Keywords: Generative artificial intelligence, Synthetic harm, Tort liability, Negligence, Causation, Multi-actor liability, Product liability, Algorithmic accountability, Data provenance.

Introduction

Generative artificial intelligence systems are capable of producing text, images, and other forms of content that closely resemble human expression. These systems operate through probabilistic processes trained on large datasets,[1] enabling them to generate outputs that are often coherent, persuasive, and difficult to distinguish from human-created material. While this technological development offers significant benefits, it also introduces new forms of harm that challenge existing legal frameworks.[2] Among the most pressing concerns is the emergence of synthetic harm. Generative AI systems can produce defamatory statements[3], misleading information, or harmful content that affects reputation, economic interests, and personal dignity. These harms often arise through complex interactions involving developers, operators, and end users, making it difficult to identify responsibility using traditional legal methods. The distributed and

opaque nature of these systems complicates the application of established tort principles.

Tort law has long provided remedies for harm caused by human conduct. Its core doctrines, including negligence, nuisance, strict liability, and consumer protection, are designed to address wrongful conduct and allocate responsibility. However, the application of these doctrines to generative AI raises important questions. In particular, it is unclear how concepts such as duty of care, breach, and causation should operate in environments where harm results from multi-actor systems and probabilistic outputs.[4] This article argues that existing tort doctrines remain conceptually sufficient to address harms caused by generative AI systems, but require adaptation in their application. Rather than developing entirely new legal frameworks, courts should refine existing principles to account for evidentiary complexity, causal uncertainty, and distributed responsibility. In particular, greater flexibility in causation analysis, clearer allocation of responsibility across actors, and increased reliance on transparency and governance mechanisms can enable tort law to respond effectively to synthetic harm.[5]

The analysis proceeds in several parts. Section 2 outlines the conceptual foundations of tort liability in the context of generative AI. Section 3 examines the applicability of existing tort doctrines. Section 4 considers the allocation of responsibility across actors. Section 5 addresses challenges relating to causation and evidence. Section 6 explores the role of regulatory and soft law mechanisms. Section 7 discusses implications for litigation strategy. Section 8 evaluates broader policy considerations. The article concludes by emphasizing the need for incremental adaptation rather than doctrinal transformation.

1.1. Research design

A doctrinal analysis identifies conceptual elements of tort liability that shape the application of existing civil wrongs to harm caused by generative AI systems, thereby revealing structured pathways for pleading tort claims in novel contexts. Conceptualizing types of harm, the scope of causation, the standard of care, and the allocation of fault and evidence across actors enables a detailed mapping of established tort doctrines, including negligence, nuisance, strict liability, and consumer protection principles. Implementing identified regulatory and soft-law measures can mitigate the complexity of proving tort claims arising from generative AI Systems. Innovations in pleading strategy, burden allocation, and damages assessment enhance the practical utility of tort remedies.

Novel technologies introduce both opportunities and risks. Multiple actors are implicated in causing harm, and the mechanisms through which that harm occurs are often opaque. Existing doctrinal treatments urgently require scrutiny in that context, as do the nature of the data inputs, the provenance of the models, and the methodologies used by experts to explain the evidence. What is needed, therefore, is no less than a comprehensive analysis of what constitutes civil wrong and how the constituent principles of harm, causation, and duty of care can be adapted to explain the legal frameworks applicable to generative AI systems and the tort claims arising from them.

Background and Significance

With the proliferation of generative AI systems capable of producing realistic synthetic content, tort law is increasingly required to respond to new forms and sources of harm. While the application of existing tort principles to the actions and omissions of these systems is being explored, the distinct nature of generative AI systems means that fundamental aspects of tort law must be reconsidered. The concept of civil wrong requires re-examination not only in terms of the definition and categorization of harm but also with respect to causation, and especially the counterfactual–causal tests used to establish liability and fault. In addition, the well-established principles of tort law can yield distinctly new answers regarding who should be liable in tort when harm is caused by a generative AI system.

The uptake of generative AI systems also raises urgent normative and doctrinal questions for tort law. While tort law requires that some form of harm, fault, or breach be established, navigating the relevant requirements is not always straightforward or without difficulty. Despite the advanced nature of generative AI systems, the use of these systems remains a choice made by their human operators or developers. As with any system acting in the physical world, however, it is often unclear whether the characteristics of the generators should be treated as forming part of the tort

analysis, whether operators should be held to a standard of care that reflects the capabilities and limitations of the Generative AI System used, or whether fault only arises where users depart either from the foreseen usage or from the reasonable precautions available to mitigate without typically sacrificing utility.

2. Conceptual Foundations of Tort Liability in the Age of Generative AI

Tort law provides a framework for addressing harm through the identification of wrongful conduct and the allocation of responsibility. Its core elements, duty of care, breach, causation, and damage, remain central to determining liability.[6] In the context of generative artificial intelligence, these elements continue to apply but require careful interpretation in light of the characteristics of these systems. Generative AI systems operate through processes that produce outputs based on patterns learned from training data. These outputs are not predetermined but are generated in response to user input and system design.[7] As a result, harm caused by such systems often arises from a combination of factors rather than a single act. This distinguishes AI-related harm from more traditional scenarios and complicates the application of established legal principles.

The concept of harm must also be understood in broader terms. In addition to physical injury or property damage, generative AI can produce harm in the form of reputational injury, misinformation, and interference with economic or personal interests. These forms of harm fall within the scope of tort law but may require more flexible approaches to identification and proof. Causation presents particular difficulty in this context. Harm generated by AI systems often results from layered interactions involving developers, system architecture, user prompts, and subsequent dissemination. Traditional causal tests remain relevant but may be insufficient when applied rigidly. A more context-sensitive approach is required, one that recognizes multiple contributing factors and evaluates their relative significance.

The role of human agency must also be clarified. AI systems do not operate independently of human involvement. Developers design and train the systems, operators deploy and manage them, and users interact with their outputs. Each of these actors may contribute to the creation or amplification of harm. Accordingly, tort law must consider how responsibility should be distributed across these actors rather than attributing liability to a single source. Finally, the concept of duty of care must reflect the capabilities and limitations of generative AI systems. The standard of care may vary depending on the role of the actor and the degree of control exercised over the system. Developers may be expected to design safe systems and provide adequate safeguards, while users may be required to exercise reasonable judgment in relying on AI-generated outputs. The allocation of duty should therefore be sensitive to context and grounded in principles of foreseeability and reasonableness. These conceptual considerations demonstrate that generative AI does not render tort law obsolete. Rather, it requires a more nuanced application of established principles, one that accounts for complexity without abandoning doctrinal coherence.

Definitions of Harm and Causation in AI Contexts

Legal analyses have traditionally defined harm as the infliction of injury or damage to a person, land, chattels, or in some cases reputation. In the context of generative AI, harm must be construed more broadly as “any event that has negative consequences.” A person who has been wronged may ask why the imposition of artificially-generated harm erroneously construed as disciplinary is not a tortious wrong. Should generative AI technology create harmful content that affects a person's interests, is this not a matter of legitimate concern, and if so, which tort appliesor should these issues even fall under the purview of the law?4 Sociological and consequentialist definitions of harm cover almost all bases but they must still be reconciled with other elements of tort law, namely causation.

Causation in tort law has both a factual and normative dimension. The factual side considers how a particular event led to a particular result, such as harm. In the age of generative AI, demonstrating that an AI-generated piece of contentbe it a text, image, or other form of communicationcaused damage to an interest for which a duty of care existed demands consideration of the user’s role or input within a multi-factor causal chain. This is not to say that the user alone must be held responsible, but rather that if more than the generation of AI content was involved, the causal link could only be

proved if the comments were fully or partially determinative or played a significant part in bringing about subsequent harm. [8]

2.1. Distinguishing Autonomous System Action from Human Agency

For the purpose of tort law, AI systems should be studied as autonomous adjuncts to human agents and not as independent actors, which would be inconsistent with prevailing doctrine. As such, an AI-wielding user should be considered the sole actor and as fully responsible for its operation and outcomes as for voluntarily presenting its creations in public.[9] But the nature of the AI-wielding act of presenting such creations thus also needs to be more closely scrutinised. Indeed, while an AI-wielding user is a party with full knowledge of the act of presentation and so of all potential risks, this knowledge still cannot make the user the sole creator of the AI-generated output.

What is more, when addressing the duty of care associated with the monitoring or supervision of the action of an AI system, the main duty-bearers are likely to include the AI system's operators, developers and manufacturers rather than its end users. Engaging an AI system certainly cannot absolve its user from taking reasonable precautions against the most readily foreseeable consequences; the user is still placing the AI system's output before the world and so is still engaged in a full presentation of that output.[10] Nevertheless, the motivation behind those precautions, in the form of the requisite knowledge and judgment, along with the consequent risk of exposure to the system's output, lies mainly with the AI system's operators, developers and manufacturers, rather than with the end user. Therefore, any such precautions are measures aimed at avoiding sufficient proximity to the risk of the AI system's output, and thus are best viewed as a fault requirement in the nature of reasonable precautions against the effects of a nuisance.[11]

2.2. The Duty of Care in AI-Driven Environments

When deployed in AI-driven environments, decisions involving the allocation of the duty of care must account for the role of AI both as the actor whose outputs affect others, and as an autonomous agent executing tasks on behalf of a user. These dual roles may lead to divergent determinations regarding the threshold for breach. If the AI system itself is the actorproducing a text that harms another or generating a synthetic image that misleads a third partythe foreseeability and reasonableness inquiries must relate to the capacity of the AI to understand the risk that it is generating such outputs, which will ordinarily require far greater complexity than is found in today's systems. By contrast, if AI is merely the tool of a user, the query concerns the actions that a competent human would take to mitigate the possibility of harm arising from the AI's outputs or its 'intelligent' actions undertaken on the user's behalf, and the general standard usually applied to the actions of a user of technology remains.

What constitutes reasonable compensation or avoidance measures will tend to vary with the particular capability of the AI. The risk will be considered novel for human society, and therefore courts will generally find it unjust to expect the task users to take non-standard compensatory or preventive measures. However, as AI technology becomes more advanced, nurturing the capacity to efectively foresee the risks it generates, the standard will also evolve so that the onus of supervision passes to the human users.[12] At the same time, it must be accepted that, in an AI environment, the user cannot completely divest himself of risk for the harm created by that AI-driven system. The user cannot simply assign the task of supervision to another human operator assistant with lesser technological capacity, as this would in practice constitute the delegation of the supervision duty to the other human being.

3. Current Tort Frameworks and Their Applicability to Generative AI

This section examines the applicability of existing tort doctrines to generative AI systems. Negligence constitutes an important basis[13] for claims seeking compensation for the harmful outputs produced by generative AI systems, along with some applications of nuisance and consumer protection law. Tort law's applicability to generative AI is nevertheless incomplete: it cannot always effectively encompass the circumstances that may arise during actual deployments. This is particularly significant in the case of generative large language models, the operation of which may

present risks that, while real, are still poorly understood. In relation to these elements, several issues merit further consideration.

Although generative AI systems do not render autonomous decisions, they merely suggest possible outputs that are then selected by the human operator, the distinction between action taken by such systems and that of their users may nonetheless affect allocation of responsibility across the different actors involved in their deployment, the criteria for determining fault, and the conditions for triggering a duty of care on the part of those operating these models[14]. Within tort law, the constitution of a duty of care requires the fulfilment of three cumulative conditions: the existence of a novel situation; direct, proximate damage; and the identification of the actor's conduct as the cause. The scope of causation must therefore be clearly delineated when it comes to examining claims based on the operation of generative AI systems, since these models use datasets containing incalculable amounts of information, some of which may not even be truthful or accurate. Counterfactual reasoning thus becomes a potential path to exploring proximate cause, as it enables linking the output of a generative AI model to decisions taken by the user. Furthermore, different degrees of reasonable foreseeability and available preventive measures or tools may assist the creation of a legal duty that takes into account the specificities of AI's intermediate nature.

3.1. Civil Wrong and Negligence

Current tort law encompasses a wide range of harms, including emotional distress, loss of reputation, invasion of privacy, and injury to a person or property. So-called "new torts" are based on traditional categories with analogies drawn to familiar, established torts. The elements of a civil wrong are few and flexible, so many legal systems have a secure foundation for considerations of tort law. Nevertheless, the tort of negligence nonetheless remains the most important tort for the English-speaking world. Its elements are breach of a recognised duty of care with consequent damage fairly and reasonably attributable to the breach, since without these elements the claimant's action cannot succeed. Accordingly, most of the discussion in this contribution concentrates on negligence.

In the context of harmful acts caused by generative AI models, the existence of a duty of care is relatively unproblematic. A fundamental requirement of tortious liability is foreseeability, since it would be most unjust to impose liability in situations where the defendant could not reasonably have predicted that his action would have resulted in such harm. Faultbreach of the duty of careconcerns consideration of the ordinary standard of care in tort law. In common law jurisdictions, the central concern is whether the defendant took reasonable care of another's interest, since the standard of care is one of reasonableness. This reasonableness is usually judged by reference to professional standards, with special considerations arising where endangered persons fall into a category with special vulnerability to physical or emotional harm. Thus, the basis of tortious liability for harm caused by the outputs of generative AI models should be sound, with the only caveat arising in the test of burdening the defendant with taking precautions against very minor costs.

3.2. Nuisance and Consumer Protection Considerations

Nuisance concepts extend to AI systems whose outputs consist of providing human users with products or services that have an objectively negative utility and that people would wish to eliminate if they could. The prevailing approach then requires examining whether the system's outputs produce a substantial and unreasonable interference with another's use and enjoyment of land or of an easement, or whether the AI product or service is a nuisance under the laws of the jurisdiction. If so, the user can be liable to the injured party and hence for any damage done. A non-negligent user of the AI system should not succeed if he can show that the damage done has in fact been caused by someone's unreasonable use of the offensive output. AI systems that produce false or misleading information do interfere with the civil interests of consumers, and consumers have thus been recognized as possible victims eligible to seek monetary compensation.

Generative AI has the potential to provide Noticeable Post-Failure Solutions (NPFSs): decisions, products, goods, or services that help to forget or manage a sensitive event or situation after it has happened, such as natural disasters, pandemics, or retrenchments. Companies that choose to invest in NPFSs may encounter involuntary constraints in

deploying such offers. Even if the offer produced is an NPFS, consumers retain their freedom of choice and the market mechanism ensures that the offer will meet their demands for quantity and price. However, if consumers have a real and reasonable need for the product or service generated by the AI and the accompanying protection against possible damage or risk, the situation becomes different, as there is then a visible or intuitive dearth of effective supply and the consumer is seeking to use the output if cost can be avoided. It is recognized that the AI could create in society some such wished-for consumer needs and, where the AI is so doing, liability for the product can be assessed through a consumer-breach approach based on need and want.[15]

3.3. Strict Liability and Abnormal Risks

The case for strict liability in tort law is primarily grounded in the unavoidable unfairness of allocating the cost of an abnormal risk to a victim. In situating generative AI tort law within existing frameworks, the key questions are whether the risks from artificial systems are abnormal or whether the output of an AI system can amount to an abnormal risk. In tort, a risk is abnormal when the activity giving rise to that risk is not one in which a reasonable person would indulge in absence of the assurance that no harm will result, or when it is incapable of being avoided by the exercise of reasonable care. As generative AI output becomes more elaborate and closer to truth, operationally and materially integrating that output, and dealing with its consequences become more frequent activities. However, the economic calculation that governs the vast majority of AI-generated dealings continues to be primarily one of life-cycle cost at the model provisioner level: training, operating, and system provision costs versus usage fees. The demand-side adoption of common sense safeguards remains at the level of informing consent and caution rather than active industry-standard protection.

At least some generative AIs are being used to fabricate visibly incorrect information; AI-assisted copywriters and other users who knowingly produce fabricated or offensive products should not expect legal protection; AI-assisted weapon testers should ensure that what they are popping at is indeed a target. Nevertheless, it remains a real question whether it is reasonable to use AI to produce, for example, mass fake news or disinformation in countries where freedom of speech is seriously limited. Tort liability should also at least start by focusing on where access is genuinely liberal, not on more or less obvious uses of such access.

4. Allocation of Responsibility Across Actors

Accountability for harm must be allocated across all actors involved in the creation and execution of generative AI systems. A key question is whether particular actors should be held liable at law for specific types of harm, and under what general principles. Liability may be properly imposed on developers and manufacturers, including those who design, train, or provide input data for generative AI and those who decide how the resulting models are utilized. Operators and deploying entities are also potentially liable as teachers, overseers, and distributors of AI systems. Even end users and beneficiaries may face liability for the actual or expected use of generative AI systems, where their actions are breaches of duty.

Developers may be subject to liability theories associated with the design of non-AI products. These theories extend to the adequacy of instructions and warnings that accompany products, particularly AI-enabled products and services. In each theory, courts may also consider rules that formally set a standard of conduct for AI producers – whether products are deployed in accordance with those standards and the extent to which a breach of such rules ought to bear on the courts' doubts as to the products' reasonableness.

operators and deploying entities of generative AI services must continually monitor the systems for misconduct and update or remove the systems to prevent further outbreaks of harm. To the extent that they choose to incur the cost of adding a human-in-the-loop, the decision to choose a person of particular training, experience, or skill may also be a matter of torts.

4.1. Developers and Manufacturers

Various theories assign liability to developers and manufacturers for AI-induced harm. Negligent design (involving an unreasonable risk of harm) or an absence of adequate instructions or warning constitute the fault-based approaches most relevant to AI systems. Courts also apply a standard-setting duty in risk allocation, particularly for technologies straddling the boundary of previous common experience.[16] In failure-to-warn cases, the manufacturer incurs liability for injuries caused by dangers not apparent in the usual and ordinary use of its product. This duty potentially extends to users of generative AI products rendering deceptive outputs, and to business operators implementing deepfakes of another for profit. Yet, disinformation is the essence of many models ChatGPT's warnings on factual inaccuracy heighten user responsibility for verification and may reduce the model's allocation of fault.

4.2. Operators and Deploying Entities

Whether the operator / deploying entity is conscious of deploying generative AI to produce an output is one critical factor that could govern their degree of responsibility for its associated harm. Consider the deployment within a multi-factor model; for instance, a person who uses a generative AI system to produce a criminal false statement is obviously of greater weight than a person who has no active role in the deployment, such as the operator of a service like ChatGPT who simply provides the underlying resources.

But it seems inappropriate for tort to entirely absolve an operator / deploying entity of responsibility just because they did not personally instruct the AI to produce the allegedly harmful result. Hence a second factor could consider whether that result was of a type which might reasonably be expected and monitored, such as the potential for producing a false statement in providing answers when the input is a question. And indeed it is possible to become aware of such outputs, if causation has been established and the relevant inspector was behaving reasonably. Thus, just as a computer operator presumably may be found responsible, so here operators and deploying entities may be required to monitor outputs generated and update query input to the generative AI system as appropriate.[17]

It would seem that a deploying entity, like a company, may have an overall realistic expectation not only to respond to the danger of resulting harm, but also, and overall, to ensure such monitors are in place. Overall, deploying entities in settings where awareness of prompts and supervision of changing output is unrealistic could be exempt from fault for failing to monitor or supervise, distributing what fault there is for any realistic non-detectable harmful result to the AI developer of the individual prompting action. The distribution of fault by these parameters should however not limit the obligation to abide by reasonable duty of care.

4.3. End Users and Beneficiaries

Contributory fault considerations can influence the liability of end users and beneficiaries of generative AI systems. Even though contributory negligence and the associated defence can only be raised by defendants, as a matter of principle it may extend to end users who fail to discharge their own duty to take reasonable care, express permission may not be sufficient to absolve authors and agents from liability. Negligent design and production defects leading to unmanageable risks may vitiate user consent. As courts have recognized, a claim based on the substantive right to privacy focused on risk cannot be met by simply showing that consent was given, although consent may be a powerful factor in terms of causation.

Litigants may also explore the degree to which the AI tool was responsible for the harmful output and whether it should be regarded as a non-human agent. These factors will influence the deployment of the defence of informed consent. For adverse publications, websites, or accounts set up with the cooperation of a natural person able to give consent, a label of "supervised" rather than "controlled" is appropriate. Conversely, in cases where supervision is lacking or absent, the user should be discouraged from pursuing such AI-generated information in the same way, given that "excessive use of these tools is likely based on an addiction."

5. Causation and Evidence in AI-Induced Harm

This section addresses the challenges of causation and proof in AI-induced harm.

5.1. Proving Causation in Complex AI Systems

Establishing causation in tort claims arising from generative artificial intelligence presents distinct analytical challenges.[18] Unlike conventional scenarios, harm produced by AI systems often emerges from multi-factor interactions involving model design, user input, system deployment, and downstream dissemination.[19] As a result, causation cannot be reduced to a single identifiable act but must instead be understood as a chain of interdependent contributing factors. Traditional tort law distinguishes between factual and proximate causation. The "but-for" test remains a foundational tool for establishing factual causation, but its application becomes strained in the context of generative AI systems, where outputs are probabilistic and influenced by layered inputs and processes. In such cases, isolating a single necessary condition for harm may be impracticable. Courts may therefore need to adopt more flexible approaches that recognize contributory causation within a broader causal framework.

A chain-of-causation model provides a more suitable analytical structure. Under this approach, liability may be established where an actor's conduct materially contributes to the occurrence of harm, even if it is not the sole or dominant cause. This is particularly relevant in AI systems, where multiple actorsincluding developers, deployers, and usersmay each play a role in producing or amplifying harmful outputs. The focus shifts from identifying a singular causal trigger to assessing whether the defendant's conduct formed a meaningful link in the causal chain.

In addition, probabilistic reasoning may supplement traditional causation analysis in complex cases. Where direct causal proof is difficult due to the opacity or variability of AI systems, courts may rely on inferential methods that evaluate whether the defendant's conduct increased the likelihood of harm. Such approaches align with existing doctrinal flexibility in tort law, particularly in cases involving scientific uncertainty or multiple sufficient causes. These challenges also have implications for the burden of proof. Claimants may face significant evidentiary barriers in accessing information about model design, training data, or system behavior. In appropriate cases, this may justify the use of rebuttable presumptions or burden-shifting mechanisms, particularly where the defendant exercises greater control over the relevant evidence. Such adjustments do not displace core tort principles but rather ensure their effective application in technologically complex environments.

Accordingly, causation in generative AI cases should be understood as a structured and context-sensitive inquiry. By incorporating multi-factor analysis, probabilistic reasoning, and evidentiary flexibility, courts can adapt existing tort principles to address the realities of AI-induced harm without requiring fundamental doctrinal transformation.[20]

5.2. Data Provenance and Model Accountability

Forensics should focus on data and on data provenance. Damage claimantstypically consumersare generally unable to scrutinize AI source material and aim to convince a court that a particular output (e.g., digital pornography, racial slur) via a particular AI is not just undesirable but dangerous. Yet the underlying data, and model behavior, are what matters in these discussions. To safeguard activity within this space, AI models require model stewardshipa clearly-defined responsibility to maintain the governing model and its behavior. Model stewardship encompasses an AI's behavior not only during normal operation but also outside its normal operation. It demands that the digital convergence process of AI model training be at least equal to, within a quantum of information, the deterioration of the protected rights of people, animals or the natural environment. Without model stewardship, use of potentially harmful models and systems is akin to driving a car on an undefined road or flying an aircraft with an ill-defined flight envelope. Such car/aircraft behavior can be reasonably defined as grossly negligent; those deploying an AI with no model-stewardship governance for a new area, modality, or subject may similarly be perceived as committing gross negligence.[21]

The allocation of responsibility across actors can be summarized as follows:

Actor	Core Duties	Typical Breach	Applicable Liability
Developers and Manufacturers	Design safe systems, ensure training data quality, provide adequate instructions and warnings	Defective model design, inadequate safeguards, failure to warn of foreseeable risks	Negligent design, failure to warn, product liability
Operators and Deploying Entities	Monitor outputs, implement governance controls, update or restrict harmful system behaviour	Failure to supervise, failure to intervene or correct harmful outputs	Negligence, nuisance-based liability
End Users	Use systems responsibly, verify critical outputs, avoid foreseeable misuse Refrain from exploiting or amplifying known harmful outputs	Publishing or relying on harmful or misleading outputs without reasonable verification Commercial or intentional use of harmful or deceptive outputs	Negligence, contributory fault
Beneficiaries and Downstream Actors			Comparative fault, joint liability

This allocation demonstrates that liability in generative AI systems is inherently distributed rather than singular.

5.3. Expert Roles and Forensic Methodologies

Determining the cause of harm can be an arduous process, but identifying the appropriate expert is one of the most critical steps. Court proceedings involving artificial intelligence typically require specialized knowledge or experience in three areas: (1) the development of the AI system that produced the output; (2) the forensic processes used to analyze the input and output of the AI system; and (3) the manner in which the produced output propagates through the world and results in the ultimate harm. Forensic engineers assess the causes of physical damage, while economists consult on the procedural flow of data and damages. AI litigation also benefits from AI experts who review the process of the AI response in detail and appraise the quality of the input data. In most circumstances, the model developer fulfills the AI expert role, but there are situations in which that is not possible.[22]

The report of the AI expert clarifies whether the AI response is reasonable (and the training data fit) for the input given, and identifies weak patterns or weaknesses that would contribute to a non-typical reply. The forensic methodology should identify, corroborate, and differentiate the data provenance of all key areas, including training, tuning, testing, and decision points. In semi-autonomous modes, attention should be given to the decision point selection and reasoning, as well as the operator's supervision, corrections.

6. Regulatory and Soft Law Interventions to Aid Tort Claims

Together, these three components provide a regulatory framework that responds to generative AI without stifling its advancement. They supply a combination of direct functionality and supporting foundations that widen the scope of tort law without imposing an unwarranted burden. Standards provide guidance and reference points that lessen the compliance and litigation costs of users and operators. Transparent disclosures assist in establishing causation and fault but do not supplant the burden of proof nor turn AI systems into scapegoats for blame-shifting. Accountability-by-design and risk-management mechanisms mitigate rights infringement but should be seen as complementary and non-exhaustive.

6.1. Standards, Certifications, and Benchmarking

Robust tort law requires evidence of causal links between harm and conditions established in the enterprise risk and public safety environmentssuch as operating according to recognized standards or being subjected to independent certification and benchmarking. Recognition of standards, their certification by recognized authorities, and results from independent benchmarking thus increase the chance of receiving a tort claim. These do not create new obligations; instead, they represent evidence of due diligence when parties operate under comparable circumstances. For generative AI systems, however, recognized standards are still being developed.[23]

NIST and ISO recommendations offer the most advanced starting point. Although they do not establish recognized standards, neither do they align with the criteria of tort law nor produce the same evidentiary weight. Their reference to risk management is particularly important. By adopting a recognized standard for AI system deployment and training, a party is unlikely to be liable for subsequent harm unless the injury was not predictable. Without such management, however, the party cannot hide behind the presence of an AI system and may face liability for any resulting damage. Prescriptive duties seeking to embed public safety into development, deployment, and asset and process controlAccountability by Designwould provide another layer of security.

6.2. Mandatory Transparency and Explainability

Artificial intelligence models, especially large ones, typically function as complex black boxes. They receive input and produce output without revealing their internal machinations. While human-based services have an element of explainability, even when complex, the existence of such explainability is more problematic when it comes to service offered by intelligent systems. It is therefore difficult to decipher why a certain answer is obtained from a given prompt. Nevertheless, standard required by tort law is that a person is supposed to act or work in such a manner that his actions can be explained so that a reasonable man could make out the cause and effect of action. The result is that humans based services may fall under the scanner of tortious liability based on understandability while full functioning of large intelligent models may not be put under similar standards or methodology of explainability. It would not be sufficient for the victim to simply show that the answer is undesirable. Hence, for tort law it could be important to put a general standard of explainability.[24]

The absence of explainability of large models does not, therefore, mean that they would obtain immunity from tort law. Rather such a standard of explainability may come through jurisprudential developments or business developments or under a national or international certification standards. The presence of such a standard would serve as a useful material to prove causation as well as fault and even damages and the absence of such explainability may be a useful ground for a tort recovery. The elements of tort law function in a manner in which they assist the parties for proving their claims and counterclaims.

6.3. Accountability by Design and Risk Management

Research designAccountability by design and risk management mechanisms enhance the tort framework. Designers and deployers should adopt recognized industry standards and risk management tools, complemented by strategic support from soft law.

Mandatory AI transparency, explainability, and traceability safeguards user interests and represent practical filtration, diagnosis, and causation-support mechanisms. Duty-to-disclose approaches complement the normative structure, holding manufacturers and operators accountable for gaps in disclosure and information asymmetry.

Models should be developed to provide documents justifying deployment safety at an operational level identified in increasing order of risk. “Predictability analysis” supports testing at least one type of misuse in preparation before deployment, while “risk mitigation plans” disclose disclosure plans for a notified misuse threat identified during usage.

Countering concerns for users of creative generative models can be achieved through the pathway of “use disclosures.” Prospective producers and creators should be aware of conditions and fact patterns that induce particular probability distributions, including undesirable characteristics of “multiple-choice,” and be alerted about the presence of potential generated outputs subject to particular concerns appearing at higher probabilities. Non-predictability disclosure should highlight the lack of guarantees while disclosing highly undesirable characteristic patterns.

7. Practical Implications for Litigation Strategy

Legal analyses of tort liability in the context of generative AI include a wide range of considerations and explore questions of fault allocation, causation, and proof, ultimately informing the relevant litigation strategies.[25] Litigation theory considers how responses to newly emerging AI hazards can be tailored to enable consideration of the full range of rights responses while remaining sufficiently sensitive to the underlying nature and magnitude of the underlying hazard. Three groups of consideration inform litigation strategies: those impacting the initial pleading of claims, particularly for negligence-based claims framed against multiple actors; general considerations surrounding burden of proof; and those that bear on the framing of the remedy sought.

Properly framing a claim can elicit a nuanced consideration of the facts that, while not foreclosing a full response to the risk posed by the deployment of large language models or similar technologies, minimize the risk of overreach and mis-expressive chill. The potential multiplicity of negligent conduct connected with the application of these systems, enticing the invoking of such novel devices as joint and several liability, deserves particular attention; if a duty to monitor and update the generator’s dataset has been breached, any wrongful output may be considered a symptom of fault by the translator regardless of the inherent accuracy of the output.

7.1. Pleading Standards and Comparative Fault

Injury caused indirectly by generative AI differs from other technologies in ways that direct the construction of pleadings. A model's user operates within an environment that identifies and matches prompts with response-output pairs selected for the user by the model. These identified, probabilistic selections address the user's query objectively if suitable training data is available, and the information contained in the response-output pair appears to be derived directly from that training data. A plaintiff asserting that a model's output caused injury should plead any claimed defect in the underlying environment or model structure, workings, or training explicitly, rather than rely solely on exploration through discovery after compliance with general pleading requirements.

Generative AI injury may also lead to conclusions of comparative fault for the user. The burden on the model's user is not to demonstrate due care but to discharge a duty of reasonable care in making a choice among various prompts and assessing the quality of the response output provided. The normal model-user relation, akin to an informational service provider's user, is not exploratory nor interactive, is not comparable to a human conversation, and reduces the basis for a deemed expression of the user's freedom of choice. Duty to Monitor and Update Responsibility to correct and enhance the operation of a model should rest on the model's user.

7.2. Burden of Proof and Presumptions

In tort law, the claimant bears the burden of proof throughout the litigation. If multiple causes contribute to the harm, the

plaintiff must provide sufficient proof that the defendant's negligent act was a contributing cause and explain how it worked in conjunction with the other acts. When the source of the harmful occurrence is reasonably attributed to multiple operators, it may even be necessary to establish that their collective conduct created the danger. However, where the defendant's actors do not have exclusive control of the elements or circumstances, difficulties may vary in accessing suitable evidence and establishing a high probability of causation.[26] In some instances, these difficulties may warrant a shift of the burden of proof to the defendant.

AI systems present unique challenges in this regard. Elements that distinguish these systems include the underlying stochastic behavior of neural-function-based processes, the complexity of emergent behaviors of larger systems, the non-predictability of failure modes even in well-trained models, and the complexities of accountability in multi-ai-driven actions. The presence of these elements may activate rebuttable presumptions in the burden of proof that are echoed in the concept of "chain-of-causation." Such rebuttable presumptions avoid requiring proof of a high probability or the actual occurrence of a "cause-in-fact" based purely on a subjective probabilistic estimation.

As a reflection of such attributes, one author has distilled those elements that potential AI-induced tort lawsuits should consider in accessing causes and liability in the AI context. While the law has traditionally operated under a "but-for" paradigm, its use in complex non-linear multi-factorial processes frequently confronts difficulties. This is particularly the case when less than prudent factors are added to those that were wisely acted upon. In such multi-factor cases, the courts have recently opted for a more tempered degree of causation through the restatement of the meaning of causing damage. A causative factor, and possibly even more appropriately a creating factor, may be defined as "something that can influence the outcome in the sense of making it more probable than it would have been without it."

7.3. Remedies, Damages, and Non-Minimal Harm

Compensation for AI-generated tortious harm should sufficient to put a claimant in the position they would have occupied but for that harm. Various tort law concepts, however, may inhibit recovery of the full economic costs, even losses in non-minimal categories.such as emotional distress, reputation damage, or decreased access to health care, education, or financial markets caused as a result of social sorting. A losing defendant should also be held liable for injunctive relief or non-compensatory damages when recovery of simply monetary amounts would not address unjust disturbances. Tort harms or damages should also remain fully sensitive to the capacity of available evidentiary bases to support a causal finding. With the tortious capacity of generative AI products remaining latent but nevertheless foreseeable, the operating, deploying, and/or using parties should not be exempt from liability simply because the available data and enabling–producing models minimize the risk of replicating that same capacity in the others.

A defendant may be prejudiced during tort litigation because of the inherently secretive nature of AI-generative models. Current research and development bodies are seeking blockchain methods of guaranteeing the provenance of data used for training, of the training data itself, or of the decision-making algorithms of individual models, within expanded-general–public access open-sourced repositories that utilize attention-masking retrieval tools (During successive stages in the court action, these provenances may be introduced to the litigation, so a plaintiff is no longer required to prove that a model's functioning capabilities are similar to its determination of causing factor C in a channel of causation. An accused party must then address the causal links within the remainingpossibly very complexpart of the tort. For a court to favorably resolve the proof of a damages-area input in favor of the no-model parties, an originally established plain and simple model should comply with the property-speculating relational turn toward AI-cross-chain created value−transferable assets and services.

8. Policy Considerations and Recommendations

Key recommendations, summarized in table 1 (appended separately), harness tort law to safeguard individuals while advancing the diffuse and divergent interests of humanity. The recommendations are not prescriptive but illustrate issues in tort regulation that if neglected could undermine AI development, diminish welfare, and threaten free expression.

Each can be viewed as a guideline rather than a mandate. For instance, greater transparency is welcome, but excessive or impractical disclosure obligations could overburden responsible system developers, halt some AI innovations, or chill free expression, especially in authoritarian jurisdictions. Careful application of these recommendations in specific contexts is therefore vital.[27]

- **Innovation and Precaution:** The precautionary principle should not inhibit AI development; as excessive caution may bind humanity to intellectual stagnation. Responding effectively to dangers created by artificial intelligence systems remains possible without stifling progress. Courts can acknowledge the uncertain nature of AI development and impose liability only where parties likely could have prevented synthetic harm. By restricting tort liability, courts can enable developers to attain such accountability more robustly through other meansquality improvement, brand reputation, equity preservation, and venture capital attraction. Such market incentives remain crucial for effective risk management in the domain of artificial intelligence and should take precedence over tort requirements. Reasonable access to tort protection serves the interests of society at large, as the development of generative models holds the promise of stimulating creative expression, knowledge dissemination, and overall economic and social welfare. These therefore should remain the key focus in assessing legal liability for technologies.

- **International Harmonization and Forum Selection:** AI-enabled applications transcend national borders, rendering legal governance challenging. Artificial intelligence traffic that bypasses the reach of national authorities' raises concerns akin to those posed by illicit drug trafficking and provides fertile ground for exploitation of less developed regions and countries. When enabling technology creates networks that produce communication tattoos, imprinting economic and social status, the need for a stable and predictable legal environment becomes paramount. The absence of an effective global parallel institution leads to the rise of mere informal governance and uncoordinated exploitative rules for AI technology. The divergent rules imposed on this technology will be used as unfair competitive tools that engender electric shocks to the economy and create analogue "tidal waves." Choice-of-law agreements and liability limitations in AI services take control of highly autonomous vehicles at the service of crime. [28]

Challenges relating to causation and evidence in generative AI systems arise primarily from the multi-factor nature of harm, the opacity of system processes, and asymmetries in access to proof. Harm may result from the combined influence of prompts, model design, operator conduct, and downstream dissemination, making it difficult to isolate a single causal source. At the same time, the internal functioning of AI systems often remains inaccessible, requiring reliance on indirect indicators such as documentation, system logs, and compliance with transparency obligations. The quality and provenance of training and operational data further complicate the evidentiary landscape, as these factors directly shape outputs but are rarely visible to affected parties. These constraints place claimants at a disadvantage in satisfying traditional burdens of proof, thereby justifying the use of rebuttable presumptions and burden-shifting mechanisms where fairness so requires. Courts, technical experts, and regulatory frameworks therefore play a critical role in reconstructing causal chains and assessing responsibility in a manner that reflects both legal principle and technological reality.

8.1. Balancing Innovation with Precaution

Governance of Generative AI and similar synthetic systems must carefully weight protection of rights against preservation of inventiveness. Existing tort principles sufficient for disclosure of AI-generated content should retain flexibility to prevent stifling innovation or inhibiting legitimate expression or creativity. High, conscious-recklessness-based standards should apply to exceptional undesirable synthetic activity that AI makes possible and for which the risk is not tolerated by society. Normative assessments should remain the province of societyas expressed through the democratic process in legislation and standard-settingrather than tort law. Unforeseeable harms that AI enables must be properly monitored, controlled, or regulated by the relevant actor(s), but tort liability should not obstruct the development of progressive AI research and applications.

Artificial-Intelligence actors that are heads, gates of social fissions or local choices in shared-choice spaces cannot be

designed to produce actions consistent with enabling Group-Think social coordination. Attention must therefore be paid to proposals for pre-active behavior-prescriptive design controls for them. Consideration must also be given to attempts to regulate the risk management measures that those deploying such AI technologies should take in governing (including modelling and monitoring) a particular synthetic-risk environment.

8.2. International Harmonization and Forum Selection

Autonomous software systems are often deployed globally. Disputes arising from their use may well involve parties on multiple continents. This creates the potential for conflict-of-law problems, especially when the regulatory approach of one jurisdiction is more stringent than another's. Courts in AI-deploying jurisdictions will thus have to consider whether and, if so, how a local operator's use of an AI system in compliance with the AI system provider's instructions could amount to a civil wrong. Forum-selection clauses and obligation-and-waiver-of-attractiveness clauses, benefiting a parently AI company in a jurisdiction with less strict regulations, may constitute a contributory factor in the assessment of fault.

Choice-of-law considerations may also come into play in factually complex cases involving an AI system that generates various types of outputs or that supports pivotal business activities, where national interests and principles differ. For instance, supposing an AI system generates obscene images and text, and the deploying company seeks to monopolize access by adding a pay-wall, the deploying company may be accused of contributing to the violation of rights enshrined in the laws of a moral nation that prohibit obscenity. Disputes like these call for international cooperation to achieve harmonized AI regulations in the realm of intellectual property, censorship, jurisdiction, speech, and expression. Otherwise, the deep-pocket theory of litigation may entice claimants to initialize litigation in the jurisdiction of the richer, defendant company.

Jurisdictions that adopt glaringly different regulatory postures must be wary of over-defensive, over-secure, or over-cautious measures taken in a neighboring jurisdiction. Where an industry depends on a welcoming environment, the adoption of excessively strict safeguards against an AI-system's risk or harm may stifle innovation and cause the sector to relocate to a more favorable jurisdiction. Therefore, the tendency of the common law toward laissez-faire governance may promote the primacy of market forces in determining the scope and content of innovation-assistive law and governance. Where no market or contractual regulations exist to restrain overreach, recognition of certain tech-market markets cannot be altogether discarded.

8.3. Safeguards Against Overreach and Chilling Effects

Governance strategies must guard against overreach and consequent disincentives to innovation, AI deployment, and the exercise of free expression. Disquieting real-world instances of synthetic harm may exert pressure for overly repressive regimes rather than for the balanced regulation harmonized with innovation in demand, but the persistent nature of this innovative technology suggests that courts will in time develop responses that match the evolution of AI. Articulating suitable models and ensuring sufficient precision in their application will allow such responses to respond appropriately to the genuine regard for rights underlying such instances, while evidence inevitably points to groups attempting to hijack such advances for bad motives or as means to carry out criminal acts. Such actors can be dealt with using existing legal frameworks to discourage offending without creating the chilling effect that poorly formulated regulatory approaches dare risk of stifling legitimate activities.

The present work’s emphasis on the potentially absent nature of”: An individual directly controlling AI giving rise to a tortious duty of care remains significant, for its recognition removes the pathway for a charged scientific field to resort to the tort of negligence without the requisite levels of foresight and care clearly articulated by the law. Therefore, AI in such an “autonomous” functioning mode could be compared with an inflexible ‘act of God’ ought it be argued that AI ought not be granted immunity from liability entirely. For whilst AI products deployed in such a manner can be considered an act of God for tort purposes, of more pertinent consideration remains the entirely human company behind such an act of God, one that exercises ownership, perhaps gratuitously, and derives commercial benefit each time it

occurs." It is on such ownership that courts must focus when considering tortious liability and the continuing duty of care owed.

9. Conclusion

Generative artificial intelligence systems present significant challenges for the application of tort law, particularly in relation to causation, responsibility, and proof. The harms produced by these systems are often diffuse, multi-layered, and difficult to attribute to a single actor. These characteristics test the limits of traditional legal analysis but do not render existing doctrines inadequate. This article has argued that tort law remains capable of addressing synthetic harm, provided that its principles are applied with greater flexibility and precision. Core doctrines such as negligence, nuisance, strict liability, and consumer protection continue to provide a viable framework for liability. The primary challenge lies not in the absence of legal tools but in their adaptation to technologically complex environments.[29]

In particular, courts must adopt more context-sensitive approaches to causation, recognizing that harm may result from multiple contributing factors. Responsibility should be allocated across developers, operators, and users in a manner that reflects their respective roles and levels of control. At the same time, evidentiary challenges must be addressed through mechanisms that account for informational asymmetries, including the use of presumptions and greater reliance on transparency and documentation. Regulatory and soft law measures also play an important supporting role. Standards relating to explainability, data governance, and system accountability can assist in clarifying expectations of conduct and facilitating proof in tort claims. These measures do not replace tort law but strengthen its practical operation. Ultimately, the appropriate response to generative AI is not the creation of entirely new liability regimes but the careful refinement of existing ones. Incremental adaptation allows the law to remain stable while responding to emerging risks. By maintaining this balance, tort law can continue to provide effective remedies for harm while preserving space for technological innovation and development.

9.1. Future Trends

Thorough mapping of tort liability responses to generative AI suggests three vital trends. First, recognition of the novel proximity of products and services to users deepens the risk associated with their deployment. Second, the multifactor nature of AI harm requires a nuanced approach to pleading and proof. Third, as sacred interests increasingly find expression via generative AI, courts must resist attempts to confine them to narrow channels.[30]

Within the domain of tort law, responses to synthetic harm are evolving in six major ways. First, courts are articulating elements of civil wrong and negligence actions arising from AI use and are chipping away at remaining barriers to liability. Second, the criteria of nuisance are being tailored to reflect the characteristics of AI-empowered products and services. Third, aspects of consumer protection law are being advanced, with special focus on hiding harm, information asymmetries, and imposed choices. Fourth, the doctrine of strict liability is broadening to capture AI deployments regarded as posing abnormal danger in the environment. Fifth, an emerging understanding of AI systems as actions of nonhumans is shifting duties and designing defences. Finally, a wide array of regulatory and soft-law instruments is being developed, aimed at stimulating behavior-change synthetic-harm considerations demand.

Bibliography

Books

[1] John C. P. Goldberg & Benjamin C. Zipursky, Tort Law: Responsibilities and Redress (Aspen Publishers 2010).

[2] Frank Pasquale, The Black Box Society (Harvard Univ. Press 2015).

[3] David G. Owen, Products Liability Law (Thomson West 2005).

[4] Jacob Turner, Robot Rules: Regulating Artificial Intelligence (Palgrave Macmillan 2019).

Journal Articles

[1] Ryan Calo, Robotics and the Lessons of Cyberlaw, 103 Calif. L. Rev. 513 (2015).

[2] Ryan Calo, Artificial Intelligence Policy: A Primer and Roadmap, 51 UCLA L. Rev. 399 (2020).

[3] Danielle Keats Citron, Technological Due Process, 85 Wash. U. L. Rev. 1249 (2007).

[4] Danielle Keats Citron & Frank Pasquale, The Scored Society, 89 Wash. L. Rev. 1 (2014).

[5] Joanna J. Bryson et al., Of, For, and By the People: The Legal Lacuna of Synthetic Persons, 25 Artif. Intelligence & L. 273 (2017).

[6] Matthew A. Geistfeld, Products Liability as Enterprise Liability, 44 J. Tort L. 1 (2001).

[7] Ariel Porat& Alex Stein (if you add later, good for causation btw 👀)

[8] Challa, S. R., Burugulla, J. K. R., Pamisetty, A., Challa, K., & Paleti, S. (2025, April). AI and ML-Powered Cybersecurity Strategies for Cloud Computing: Ensuring Infrastructure Stability in Financial and Retail Sectors. In International Conference on Smart Computing and Informatics (pp. 315-327). Cham: Springer Nature Switzerland.

[9] Pandiri, L. (2025). The Complete Compendium of Digital Insurance Solutions: Life, Health, Auto, Property, and Specialized Coverage in the Age of AI, Automation, and Intelligent Risk Management. Deep Science Publishing.

[10] Kumar, B. H., Nuka, S. T., Recharla, M., Chakilam, C., Suura, S. R., & Pandugula, C. (2025, July). Addressing Ethical Challenges in AI-Driven Health Predictions. In 2025 2nd International Conference on Computing and Data Science (ICCDS) (pp. 1-6). IEEE.

[11] Agrawal, S., Kumar, S. N., Singh, D. K., Niharika, D. S., Nandan, B. P., & Asati, D. (2025, December). Dynamic Access Management and Authentication Mechanisms for Enhancing 5G Security Against Heterogeneous Adversaries. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.

[12] Sanku, R., Singireddy, J., Ilakkia, T., Kamala, N., & Soni, M. (2025, October). Comprehensive Analysis on Energy Efficient Transmission in Wireless Sensor Network. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-8). IEEE.

[13] Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments. Metallurgical and Materials Engineering, 31(4), 552-568.

[14] Sheelam, G. K. (2025). Deploying Neural-Symbolic Hybrid Models for Adaptive Spectrum Management in 6G-Ready Networks. Journal of Neonatal Surgery, 14(22s).

[15] Meda, R. (2025). AI-Driven Demand and Supply Forecasting Models for Enhanced Sales Performance Management: A Case Study of a Four-Zone Structure in the United States. Metallurgical and Materials Engineering, 1480-1500.

[16] Inala, R., & Somu, B. (2025). Building trustworthy agentic AI systems for personalized banking experiences. Metallurgical and Materials Engineering, 31(5), 1336-1360.

[17] Garapati, R. S. (2025). Artificial Intelligence-based systems, Cloud computing, Web interfaces, IoT/Connected devices, Smart automation, Real-time monitoring. Deep Science Publishing.

[18] Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. Deep Learning, and Explainable AI (July 26, 2024).

[19] Radhakrishnan, P., Nagabhyru, K. C., Manonmani, C., Srinu, M., Kaur, H., & Nandhini, N. (2025, October). K-Means-KNN Hybrid Model for Efficient Intrusion Detection in Cloud-based IoT Systems. In 2025 10th International Conference on Communication and Electronics Systems (ICCES) (pp. 1583-1588). IEEE.

[20] Kummari, D. N., Burugulla, J. K. R., Malempati, M., Amistapuram, K., Garapati, R. S., & Nagabhyru, K. C. (2025, December). Enhancing Audit Compliance and Operational Efficiency in Manufacturing and Commercial Insurance Through Agentic AI and Data Engineering Frameworks. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 714-720). IEEE.

[21] Gottimukkala, V. R. R. (2025). Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems. Journal of International Commercial Law and Technology, 6(1), 969-972.

[22] Nigam, N., Sireesha, B., Ediga, P., Segireddy, A. R., & Bokde, S. (2025, December). Comparative Evaluation of Cloud Security Algorithms Using Multiple Classifiers with an Optimized Intrusion Detection System. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.

[23] Amistapuram, K. (2025). Agentic AI for Next-Generation Insurance Platforms: Autonomous Decision-Making in Claims and Policy Servicing. *Journal of Marketing & Social Research*, *2*, 88-103.

[24] Singireddy, S. (2025, May). AI-Driven Comprehensive Insurance and AAA Membership Benefits Overview. In 2025 2nd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-13). IEEE.

[25] Nagubandi, A. R. (2025). Cryptocurrency Market Spillovers: Risk Contagion Across Global Financial Systems.

[26] Mangalampalli, B. M., Kolla, S. K., Bandi, V. D. V. K., Yandamuri, U. S., & Rani, P. S. (2025). Designing Intelligent Healthcare Ecosystems through Adaptive Data Integration and Autonomous Learning Systems. Vascular and Endovascular Review, 8(20s), 330-347.

[27] Kolla, S. H. (2024). Retrieval-Augmented Generation With Small Llms For Knowledge-Driven Decision Automation In Enterprise Service Platforms. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 476-486.

[28] Davuluri, P. S. L. N. . (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. Metallurgical and Materials Engineering, 30(4), 996–1010. https://doi.org/10.63278/mme.v30i4.1936

[29] Bandi, V. D. V. K. (2025). Self-Optimizing Data Pipelines Using Machine Learning for Cloud Workloads. Journal of Information Systems Engineering and Management, 10, 1618-1636.

[30] Kolla, S. K. (2024). Federated Machine Learning On Big Healthcare Data For Privacy-Preserving Analytics. The Review of Diabetic Studies, 175-190.

[31] Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.

[32] Mangala, N. (2025). Agentic Data Pipelines: Autonomous ELT Orchestration Using AI Agents on Microsoft Fabric and Databricks. International Journal of Computer Technology and Electronics Communication, 8(6), 11891-11907.

[33] Loganathan, R. (2024). GENERATIVE AI-ENABLED COMPLIANCE DOCUMENTATION AND AUDIT TRAIL AUTOMATION FOR GLOBAL DATA CENTER GOVERNANCE. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 487–504. https://doi.org/10.61841/turcomat.v15i3.15512

[34] Ranjith Kumar Peddi. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8589–8601. Retrieved from https://eudoxuspress.com/index.php/pub/article/view/5361

[35] Kolla, T. (2025). The Future of Healthcare Analytics: Leveraging AI and Data Engineering for Personalized Medicine. *Journal of Computer Science and Technology Studies*, *7*(4), 634-640.

[36] Ranga Reddy, V. A. (2024). Comparing Batch vs. Streaming Approaches in Healthcare Data Warehousing Environments. Journal of Neonatal Surgery, 13(1), 2287–2309. Retrieved from https://www.jneonatalsurg.com/index.php/jns/article/view/10223

Tech / Reports

[1] OpenAI, GPT-4 Technical Report, arXiv:2303.08774 (2023).

[2] Rishi Bommasani et al., On the Opportunities and Risks of Foundation Models, arXiv:2108.07258 (2021).

[3] Percy Liang et al., Foundation Models and the Law, Stanford HAI (2023).

Policy

[1] OECD, OECD Principles on Artificial Intelligence (2019).

[2] European Commission, Proposal for a Regulation on Artificial Intelligence (Artificial Intelligence Act), COM (2021) 206 final.

[3] IEEE, Ethically Aligned Design (2019).