AI-Assisted Decision-Making and Emerging Tort Liability

BY RICK WAGNER

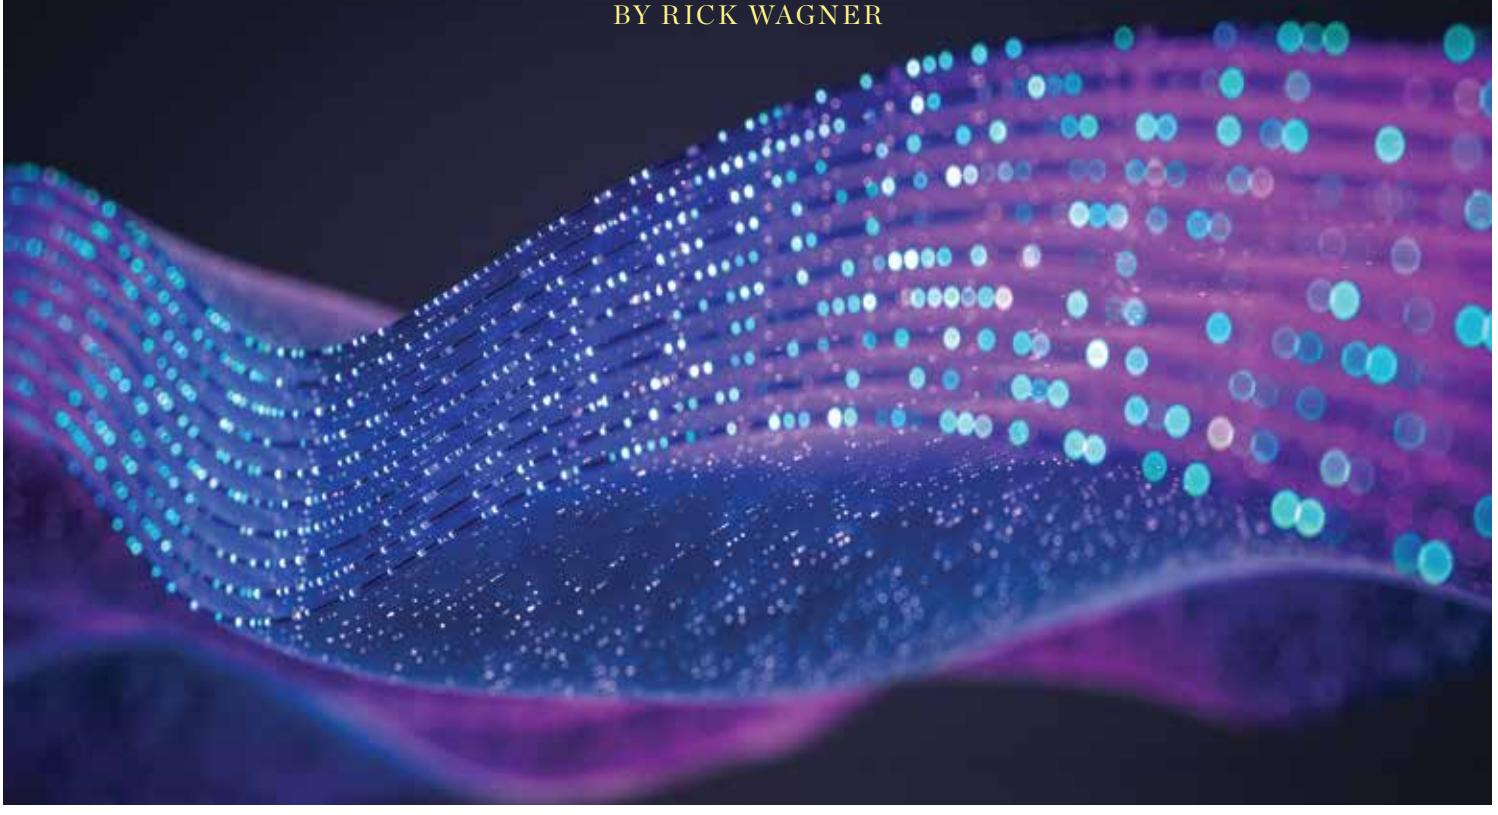

As AI-assisted systems increasingly influence institutional decisions, traditional negligence principles—duty, foreseeability, and causation—are likely to govern liability. Earlier treatment of algorithmic tools provides guidance for emerging AI-related claims.

Artificial intelligence and algorithmic tools are increasingly used to assist decision-making in fields ranging from insurance claims handling to medical triage and risk assessment. Systems using these tools are commonly described as "advisory," intended to support—rather than replace—human judgment. When harm occurs, defendants may contend that responsibility is diluted because a machine generated the recommendation that shaped the outcome.

From a tort perspective, however, the presence of an AI-assisted tool does not alter the basic analytical framework. Courts evaluating negligence claims have long relied on familiar principles: duty, breach, foreseeability, causation, and damages. Colorado negligence law, in particular, frames liability around those familiar elements.[1] What has changed is not the law, but the architecture of decision-making. The central question is no longer whether a system was used, but how it was used—and whether decision authority was structured in a way that made resulting harm foreseeable.

Existing Tort Principles Still Govern

Courts assessing negligence claims focus on whether a defendant owed a duty of care, whether the risk of harm was foreseeable, and whether the defendant's conduct was a cause of the plaintiff's injury.[2] In Colorado, a duty to use due care generally "arises in response to a foreseeable and unreasonable risk of harm to others."[3] Foreseeability is a practical concept. It asks whether the harm that occurs falls within the scope of danger created by the defendant's conduct.[4]

Defendants have long attempted to soften liability by pointing to intervening processes: internal guidelines, third-party vendors, standardized decision trees, or other "systems" said to interrupt the causal chain. But it is well-settled that even "an intentionally tortious or criminal act of a third party is not a superseding cause immunizing [a] defendant from liability . . . ."[5] Thus, where an organization adopts a decision structure, fixes its parameters, and builds that structure into daily workflow, the system is not an independent actor. It is simply how the organization chooses to act, whatever labels appear in a policy manual. Responsibility follows authority. If the organization retains authority over adoption, configuration, incentives, and use, it remains responsible for foreseeable harm that flows from that structure.

AI-assisted systems fit comfortably within this framework. The question does not turn on whether a tool is intelligent, adaptive, or novel, but on whether its role in decision-making predictably affects outcomes in ways that create foreseeable risk.

Algorithmic Decision Systems Before AI

Courts and regulators confronted algorithmic decision-making long before the term "artificial intelligence" entered common legal usage. In the insurance context, one of the most prominent examples was Colossus, a computer program widely adopted by insurers in the late 1990s and early 2000s to evaluate bodily injury claims.

Colossus was not artificial intelligence in the modern sense. It did not learn, adapt, or operate autonomously. Instead, it applied fixed variables, valuation ranges, and weighting rules to claim inputs in order to generate settlement recommendations. Insurers consistently characterized the program as advisory, emphasizing that adjusters retained discretion to deviate from its outputs.

Litigation and regulatory investigations, however, revealed that Colossus often functioned as a decision anchor, exerting substantial pressure on settlement behavior even when nominal human review existed. Courts and litigants examining Colossus focused on how insurers integrated the program into claims-handling workflow and whether meaningful adjuster discretion remained in practice.[6]

Regulators took a similar approach. State insurance regulators investigated Allstate's use of Colossus and entered a multistate regulatory settlement addressing how the software was used in claims practices.

Taken together, Colossus-era litigation and regulation reflect a consistent treatment of algorithmic decision systems: they are analyzed as part of the organization's conduct, not as independent actors. When software meaningfully shapes outcomes, responsibility remains with the entity that selected, configured, and relied upon it.

Assistance Versus Substitution

An issue that could affect litigation involving automated systems is the distinction between assistance and substitution.

A tool may genuinely assist decision-making when it provides information that is meaningfully evaluated by a human decision-maker—one who understands the system's limits, can question its conclusions, and can override its recommendations. In that setting, the system functions much like a checklist, database query, or statistical report.

Substitution occurs when the system's output effectively determines the outcome, even if nominal human review remains. This shift rarely requires explicit policy language. More often, it emerges through workload pressure, performance metrics, or internal practices that make deviation from the system costly, slow, or professionally unrewarding. Over time, the tool becomes dispositive in practice, if not in name.

From a liability standpoint, that boundary matters. As reliance increases, so does foreseeability. When decision authority is functionally delegated to an automated system, harms that predictably flow from that delegation remain

attributable to the organization that designed and maintained the decision structure.

A Commonplace Hypothetical

Consider a hypothetical health insurance carrier that rolls out an AI-assisted claims review tool designed to flag claims for additional scrutiny based on historical data patterns. The system evaluates treatment duration, billing codes, prior claim history, and injury descriptions, producing a risk score for each claim.

Official policy describes the tool as advisory. Adjusters are instructed to review flagged claims independently before deciding whether to approve, deny, or delay payment. In real-world operation, however, adjusters carry high caseloads and are evaluated on efficiency metrics. Claims flagged by the system are routinely delayed or denied unless an adjuster creates a paper trail explaining why the system should be ignored.

A claimant with a legitimate injury experiences prolonged delay in treatment approval, resulting in worsened medical outcomes. The carrier argues that no adjuster intentionally denied care and that the AI-assisted tool merely supplied information.

Under traditional tort analysis, the focus is not on the intent of individual adjusters or the advisory label attached to the system. The relevant questions are whether the risk of harm from delayed treatment was foreseeable and whether reliance on the system was a substantial factor in producing the harm.[7] How the carrier structured use of the system may bear directly on both questions.

Courts confronting earlier algorithm-mediated systems, including Colossus-related claims handling, have treated those questions as matters of organizational responsibility rather than technological novelty.[8]

Foreseeability in Automated Contexts

Foreseeability does not require precise prediction of the way harm occurs. Colorado courts have recognized that proximate cause may exist even when an actor did not—and could not—foresee the exact mechanism of injury.[9] The inquiry is whether a reasonably careful person, under similar circumstances, would anticipate that injury might result from the conduct at issue.[10]

> “When organizations deploy systems that constrain decision-making, downstream effects are often foreseeable precisely because the systems are designed to standardize outcomes and reduce variance. Modern AI systems increase the scale, speed, and opacity of such constraints, but they do not alter the underlying inquiry.”

When organizations deploy systems that constrain decision-making, downstream effects are often foreseeable precisely because the systems are designed to standardize outcomes and reduce variance. Modern AI systems increase the scale, speed, and opacity of such constraints, but they do not alter the underlying inquiry.

When organizations deploy systems that constrain decision-making, downstream effects are often foreseeable precisely because the systems are designed to standardize outcomes and reduce variance. Modern AI systems increase the scale, speed, and opacity of such constraints, but they do not alter the underlying inquiry.

Causation and Responsibility

Defendants may argue that automated systems introduce intervening causes that break the chain of causation. Colorado law, however, treats intervening causes as part of the foreseeability analysis. An intervening force is one that actively operates in producing harm after the actor's negligent act or omission has occurred.[11] An intervening act generally does not absolve a defendant of responsibility if it is reasonably foreseeable.[12] Only where the intervening act is unforeseeable and fully independent may it constitute a superseding cause that breaks the causal chain.[13]

Where an organization selects a system, defines its parameters, and embeds it into operational decision-making, the system's outputs are not independent acts. They are part of the organization's conduct. Causation therefore turns on whether reliance on the system was a substantial factor in producing the harm.[14]

Affirmative Defenses in AI-Assisted Decision-Making Cases

Because the governing framework remains orthodox negligence, defendants predictably assert orthodox defenses. The presence of an AI-assisted tool does not eliminate comparative fault, assumption-of-risk concepts, enforceable waivers, or superseding cause arguments. What changes is how those defenses are evaluated when decision-making is mediated by software.

Colorado applies comparative negligence, reducing damages in proportion to the plaintiff's percentage of fault and barring recovery if the plaintiff's negligence equals or exceeds the combined negligence of the defendant or defendants.[15] Implied assumption of risk is treated as a form of contributory negligence and is generally subsumed within the comparative negligence framework.[16] In cases involving algorithmic or AI-assisted decision systems, defendants may argue that plaintiffs failed to provide information, delayed reporting, or otherwise contributed to the harm. Compar-

ative negligence may allocate fault, but it does not negate an organization's duty to structure and supervise decision systems in a way that does not foreseeably constrain judgment or produce systemic harm.

Defendants may also argue that plaintiffs consented—by contract, disclosure, or form—to the use of automated tools or to limitations on liability. Colorado recognizes enforceable exculpatory agreements in appropriate circumstances, but courts scrutinize them carefully and do not permit releases to shield willful and wanton conduct.[17] Where the operative mechanism of harm is not the mere use of software, but its institutional role in decision-making, courts may examine whether any purported consent meaningfully described that role or merely relied on generalized boilerplate.

Finally, defendants may attempt to characterize software outputs—or subsequent human reliance on them—as intervening causes. Colorado's proximate-cause jurisprudence largely neutralizes that move where the intervening conduct is foreseeable.[18] The central factual question is whether reliance on the tool was fully independent of the defendant's own decisions about adoption, incentives, staffing, training, and oversight, or whether it was the predictable downstream effect of the structure the defendant put in place.

Oversight and Delegation

Liability exposure frequently turns on questions of oversight. Organizations remain responsible for understanding the limits of tools they deploy, monitoring their real-world effects, and adjusting practices when unintended consequences appear. In Colorado's proximate-cause doctrine, foreseeability rests on common-sense perceptions of the risks created by conditions and circumstances.[19]

Delegation of authority does not absolve responsibility; it reallocates risk. When decision-making is structured in a way that predictably produces harm, the law looks to those who designed and maintained that structure. AI-assisted systems are no different in this respect from staffing models, scheduling practices, or incentive schemes that have long been examined under negligence principles.

Colorado's Emerging AI Regulatory Framework

Recent legislative activity underscores the relevance of these issues without altering the underlying tort analysis. In 2024, the Colorado General Assembly enacted Senate Bill 24-205, commonly referred to as the Colorado Artificial Intelligence Act, which imposes reasonable-care obligations on developers and deployers of specified high-risk AI systems and targets known or reasonably foreseeable risks of algorithmic discrimination in consequential decisions.[20]

Implementation has been deferred amid debate regarding scope and practical effect. In 2025, the legislature extended the effective date of SB 24-205's requirements to June 30, 2026.[21] Whatever the statute's ultimate contours, its emphasis on reasonable care, foreseeable risk, and oversight tracks negligence concepts courts already apply to institutional decision-making. Regulatory developments do not replace tort doctrine; they reflect parallel concern that delegating consequential decisions to automated systems carries foreseeable risks that remain the responsibility of those who deploy them.

Conclusion

AI-assisted decision-making does not create a new category of tort liability. It presents a modern variation on familiar problems: delegation of judgment, institutional reliance on standardized processes, and foreseeable harm resulting from constrained discretion.

Courts have already demonstrated, through earlier treatment of algorithmic decision systems, that established negligence principles adapt readily to new decision architectures. Modern AI expands the bandwidth, scale, and opacity of such systems, but it does not displace the legal framework that governs responsibility. Judgment cannot be outsourced without consequence. CL

Rick Wagner is a Colorado trial attorney with more than 25 years of experience in personal injury litigation. He has previously worked in law enforcement and prosecution and has handled cases involving complex questions of liability and organizational decision-making.

Coordinating Editor: Jennifer Seidman, jseidman@hearnfleener.com

NOTES

1. CJI-Civ, ch. 9 (Negligence—General Concepts).
2. *United Blood Servs. v. Quintana*, 827 P.2d 509, 519 (Colo. 1992).
3. *Id.*
4. *Taco Bell, Inc. v. Lannon*, 744 P.2d 43, 48 (Colo. 1987).
5. *Ekberg v. Greene*, 588 P.2d 375, 376–77 (Colo. 1978).
6. *See Truong v. Allstate Ins. Co.*, 227 P.3d 74 (N.M. 2010); *Kosierowski v. Allstate Ins. Co.*, 51 F.Supp. 2d 583 (E.D.Pa. 1999).
7. *Build It and They Will Drink, Inc. v. Strauch*, 253 P.3d 302, 306 (Colo. 2011).
8. *See Truong*, 277 P.3d 74; *Kosierowski*, 51 F.Supp.2d 583.
9. *Deines v. Atlas Energy Servs., LLC*, 2021 COA 24, ¶ 13.
10. *Taco Bell*, 744 P.2d at 48.
11. *Restatement (Second) of Torts* § 441(1) (American Law Institute 1965).
12. *Ekberg*, 588 P.2d at 377.
13. *Albo v. Shamrock Oil & Gas Corp.*, 415 P.2d 536, 537 (Colo. 1966).
14. *Restatement (Second) of Torts* § 431, *supra* note 11.
15. CRS § 13-21-111.
16. *Brown v. Kreuser*, 560 P.2d 105 (Colo.App. 1977).
17. *Jones v. Dressel,* 623 P.2d 370 (Colo. 1981).
18. *Deines*, 2021 COA 24, ¶ 13.
19. *Taco Bell*, 744 P.2d at 48.
20. SB 24-205, 74th Gen. Assemb., 2d Reg. Sess. (Colo. 2024).
21. SB 25B-004, 75th Gen. Assemb., Extraordinary Sess. (Colo. 2025).