Hogan Lovells

Insights and Analysis

Negligence & AI: Can the courts keep up?

06 May 2026

At this early stage, be cautious in how you talk about your commitment to AI best practices. Without a single federal standard governing what harmful use of AI looks like, courts are continuing to take up AI-related cases, establishing in real-time the bounds of AI liability. As Elizabeth A. Och of Hogan Lovells writes, two likely upshots are increased fragmentation and inconsistent approaches across jurisdictions.

Hardly a day passes without a headline about a new lawsuit tied to an AI chatbot or other AI-enabled system. In the absence of comprehensive federal legislation and a corresponding private right of action for AI harms, plaintiffs are increasingly turning to common-law tort theories to frame these claims. Negligence has emerged as an attractive vehicle for plaintiffs. It is flexible, broadly available, familiar to courts and can be asserted in virtually any court and adapted to a wide range of factual scenarios.

These cases are rarely “about” AI in the abstract. Instead, they focus on human judgment: how AI systems were designed, selected, governed, deployed, monitored and constrained, or in some cases, how they were not. As AI models grow more sophisticated and agentic systems move from

experimental deployments into everyday use, negligence claims are likely to become more frequent and more consequential.

At the same time, courts move slowly. Cases are often resolved years after the events in question, during which the underlying technology and industry practices will evolve. Courts will be asked to assess what was “reasonable” or “foreseeable” at a particular moment in time, using doctrines developed for more stable technologies. Whether courts can adapt traditional negligence principles to this rapidly changing landscape — and do so consistently — will shape the contours of AI liability in the years ahead.

Who can be sued and for what duty?

Negligence claims can target human or corporate actors across the AI lifecycle. Plaintiffs may name developers, model providers, integrators, deployers, platforms or even end users, particularly in high-stakes professional or commercial settings. Importantly, plaintiffs are not required to identify a single “responsible” actor at the outset. Instead, they can sue broadly and allow discovery to reveal where decision-making authority and **risk control** resided.

Early complaints reflect an expansive view of potential duties. Plaintiffs may allege that a defendant (depending on its position in the AI lifecycle) had obligations to design and train systems responsibly; to use appropriate and representative data; to anticipate foreseeable misuse; to warn of known or reasonably knowable limitations; to implement safeguards against **risks**; to ensure meaningful human oversight; to conduct use-case-specific testing; to train downstream users; to enforce terms of use and safety policies; to select AI tools that were appropriate for the task at hand; and to avoid blind reliance on AI outputs in contexts requiring independent judgment.

Whether any such duties exist is ultimately a question for the courts. The analysis turns on familiar negligence considerations, such as the nature of the relationship between the parties, the degree of control exercised by defendants and the context in which the conduct occurred. These inquiries often cut across corporate boundaries and contractual layers, complicating early attempts to narrow the case.

Reasonable care without a federal benchmark

The success of a negligence claim frequently hinges on whether the defendant exercised “reasonable care.” In the AI context, that inquiry is complicated by the absence of a single, authoritative federal standard of care. Courts are likely to assemble the reasonable-care benchmark from a patchwork of sources, including industry practices, internal policies, voluntary frameworks, expert testimony and post-hoc assessments of what precautions could have been taken.

Defendants may point to this regulatory vacuum as a defense, arguing that no settled industry standards existed at the relevant time, that they followed prevailing practices and that plaintiffs are attempting to impose hindsight-driven expectations. Evidence of **compliance** with existing regulatory regimes, such as consumer protection laws, professional standards or sector-specific safety requirements, may be offered as proof of reasonableness.

Plaintiffs may use the same federal statutory void as a reason to look inward. Internal AI policies, **governance** frameworks and aspirational public statements may be cited as evidence of the applicable standard of care. Language intended to signal a commitment to best practices can be reframed as a self-imposed duty that the organization failed to meet. The lesson is not that companies should avoid adopting AI policies but that such policies should be realistic, risk-tiered and demonstrably implemented rather than purely aspirational.

Known risks, evolving systems

Negligence liability extends only to harms that are foreseeable. Plaintiffs can increasingly point to widely recognized categories of AI risk — such as bias, hallucinations, data drift, misuse and overreliance — as evidence that harm of the general type was foreseeable, even if the precise outcome was not.

For courts, the central question is often not whether a particular outcome could have been predicted in advance but whether reasonable actors should have anticipated the relevant risk category and taken proportionate steps to mitigate it. As a result, documentation matters. **Risk assessments**, testing

protocols, monitoring practices and incident-response procedures may all play a critical role in evaluating foreseeability.

No negligence claim can succeed without proof of causation, and this element may prove the most challenging for plaintiffs in AI cases. Model behavior can be opaque, outputs may vary based on prompts and context, and multiple human actors often intervene between system output and the alleged harm. Model updates, retraining or version changes can further complicate efforts to identify which system caused a particular injury.

Defendants, for their part, may attempt to push liability upstream or downstream, emphasizing their own lack of control, the absence of a direct relationship with the plaintiff or intervening human decisions. Traceability becomes a strategic asset. Version control, audit trails, documentation of human review and records of incident detection and remediation can all influence whether causation arguments are resolved early or survive into costly discovery.

The role of courts & the risk of inconsistent outcomes

As AI systems continue to evolve, courts' application of negligence principles will develop alongside them. Courts draw on the existing precedent in their jurisdiction (or suitable analogies from similar cases), which may take the law in different directions depending on the jurisdiction. Some may treat downstream misuse or overreliance as an intervening cause that breaks the causal chain; others may view such conduct as foreseeable, particularly where warnings or safeguards were inadequate. What was unforeseeable one year may be deemed foreseeable the next.

Judicial philosophy and technical familiarity with AI systems will also play a role. Some courts may be reluctant to expand tort liability in the absence of legislative guidance, while others may view tort law as a necessary gap-filler. At the same time, plaintiffs are unlikely to rely on negligence alone, instead pairing it with claims **under state AI laws**, consumer protection statutes, product liability theories or other available causes of action. The result will be increased fragmentation, forum shopping and inconsistent approaches across jurisdictions.

This article was originally published for ***Corporate Compliance Insights*** *here.*

Authored by Elizabeth Alice “Liz” Och.

Contacts

Elizabeth (Liz) Alice Och

Counsel

Denver

Email me