Date of Hearing: May 6, 2025
Fiscal: No

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION
Rebecca Bauer-Kahan, Chair
AB 316 (Krell) – As Amended April 28, 2025

**SUBJECT**: Artificial intelligence: defenses

## SYNOPSIS

*The hallmark of artificial intelligence (AI) is the ability of a machine to learn and act autonomously. Autonomous systems, particularly those that use machine-learning, can behave in unpredictable ways that are not foreseeable to developers or deployers. When such systems are implicated in harm to a person or property, courts may struggle to apply traditional liability frameworks that rely on concepts such as the foreseeability of harm and the blameworthiness of a person's conduct.*

*This bill provides that in a lawsuit against a developer or user of an AI alleged to have harmed the plaintiff, the developer or user cannot assert the defense that the AI was acting autonomously. The bill is co-sponsored by Children's Advocacy Institute at the University of San Diego, School of Law and Organization for Social Media Safety and supported by several civil society organizations. Supporters argue that the bill promotes accountability by making it clear that humans or corporations who develop or use AI cannot evade liability by attributing harm to the AI itself.*

*The bill is opposed by Chamber of Progress and TechNet. They argue the bill is premature, will stifle innovation, and will lead to greater liability.*

*The Judiciary Committee passed the bill by a 12-0 vote.*

**THIS BILL**:

1) Incorporates the definition of "artificial intelligence" described below.

2) Provides that in a civil action against a defendant who developed, modified, or used AI that is alleged to have caused a harm to the plaintiff, it shall not be a defense, and the defendant may not assert, that the AI autonomously caused the harm to plaintiff.

**EXISTING LAW**:

1) Defines "artificial intelligence" as "an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments." (Civ. Code § 3110(a).)

2) Provides that every person is bound, without contract, to abstain from injuring the person or property of another, or infringing upon any of their rights. (Civ. Code § 1708.)

3) Provides that anyone willfully deceives another with intent to induce them to detrimentally alter their position is liable for damage suffered as a result. (Civ. Code § 1709.)

4) Provides that anyone who obtains a thing without consent, as specified, must restore it to the person from whom it was thus obtained, except as specified. (Civ. Code § 1712.)

5) Provides that everyone is responsible, not only for the result of their willful acts, but also for an injury occasioned by their want of ordinary care or skill in the management of their property or person, except so far as the latter has, willfully or by want of ordinary care, brought the injury upon themselves. (Civ. Code § 1714(a).)

6) Provides that manufacturers and sellers are not liable in product liability actions if the product is inherently unsafe and known to be unsafe and is a common consumer product intended to personal consumption. (Civ. Code § 1714.45(a).)

**COMMENTS**:

1) **Author's statement**. According to the author:

> The California AI industry is rapidly growing, both from an economic and technological standpoint. AI has seen extraordinary advancements in its applications, complexity, and autonomy, to the point where AI is replacing human intelligence in certain tasks. As AI becomes more complex, it is increasingly involved in daily interactions and significant decision-making. While this has the potential to bring positive changes to various industries and facets of life, this also means that AI related harm can be much more significant. These harms are already manifesting and will only worsen as the AI race becomes more competitive. Specifically, AI being deployed through social media has been shown to be particularly harmful to youth.
>
> This bill ensures that companies benefiting from the use of AI are also responsible for the harms AI may cause. By eliminating a potential AI defense theory, this bill encourages careful vetting of AI products before they are used and ensures that there is a legal entity held to account if AI is shown to violate the law.

2) **Artificial intelligence**. AI refers to the mimicking of human intelligence by artificial systems such as computers. AB 2885 (Bauer-Kahan, Stats. 2024, Ch. 843) defined the term as "an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments." AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process: numbers, text, audio, video, or movement. AI is not fundamentally different from other computer functions; unlike other computer functions, however, AI is able to accomplish tasks that are normally performed by humans.

Most modern AI tools are created through a process known as "machine learning." Machine learning involves techniques that enable AI tools to learn the relationship between inputs and

outputs without being explicitly programmed.[1] The process of exposing a naïve AI to data is known as "training." The algorithm that an AI develops during training is known as its "model." At its core, training is an optimization problem: machine learning attempts to identify model parameters – weights – that minimize the difference between predicted outcomes and actual outcomes. During training, these weights are continuously adjusted to improve the model's performance by minimizing the difference between predicted outcomes and actual outcomes. Once trained, the model can process new, never-before-seen data.

Models trained on small, specific datasets in order to make recommendations and predictions are sometimes referred to as "predictive AI." This differentiates them from generative AI (GenAI) which are trained on massive datasets in order to produce detailed text, images, audio, and video. When ChatGPT generates text in clear, concise paragraphs, it uses GenAI that is trained on the written contents of the internet.[2] When Netflix suggests content to a viewer, its recommendation is produced by predictive AI that is trained on the viewing habits of Netflix users.[3]

3) **Autonomous AI.** Unlike traditional automated systems that are pre-programmed and have a predictable narrow outcome, AI systems – particularly those that use machine learning and deep learning[4] – can produce outputs that are not pre-determined by inputs, enabling them to act in novel and unpredictable ways.[5] "This constitutes autonomous decision-making. As machine-learning and deep-learning capabilities advance, AI systems may be technically able to make predictions independently. AI systems may act in ways that humans would not have considered, reducing the control humans have over the outcomes."[6] Examples of scenarios in which AI may behave autonomously or unexpectedly follow.

*AI agents.* Advanced AI systems can use "chain of thought" problem-solving by breaking down complex tasks into a sequence of intermediate steps, making progress towards developing AI agents: "general-purpose AI which can make plans to achieve goals, adaptively perform tasks involving multiple steps and uncertain outcomes along the way, and interact with its environment – for example by creating files, taking actions on the web, or delegating tasks to other agents – with little to no human oversight."[7] AI agents have been tested, with some success, for tasks such as online shopping, assistance with scientific research, software development, training machine learning models, carrying out cyberattacks, and controlling robots. Progress in this area is rapid.[8]

---

[1] IBM, *What is machine learning?* www.ibm.com/topics/machine-learning.
[2] OpenAI, *How ChatGPT and Our Language Models Are Developed*,https://help.openai.com/en/articles/7842364-how-chatgpt-and-our-foundation-models-are-developed.
[3] Netflix, *How Netflix's Recommendations System Works*, https://help.netflix.com/en/node/100639.
[4] Deep learning is a subset of machine learning that uses multilayer neural networks to analyze vast datasets. (Amazon Web Services, "What is Deep Learning?" https://aws.amazon.com/what-is/deep-learning/.)
[5] Buiten et al, "The law and economics of AI liability," 48 Computer Law & Security Review (Apr. 2023), https://www.sciencedirect.com/science/article/pii/S0267364923000055#cit_61.
[6] *Ibid.*
[7] "International AI Safety Report," AI Action Summit (Jan. 2025), p. 38, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.
[8] *Id.* at p. 44.

*Sorcerer's Apprentice.* Models that use reinforcement learning – a training process that uses rewards and punishments to orient a model's behavior towards a specific goal[9] – can sometimes attain the goal in unexpected ways. Dario Amodei, co-founder and CEO of Anthropic, famously experienced this when he was developing an autonomous system that taught itself to play a boat-racing video game. The system discovered that it could maximize its goal of scoring points not by winning the race but by driving in circles, colliding with other boats, and catching on fire inside of a harbor with replenishing power-ups that allowed the system to rack up more points than simply winning the race.[10] Like in Johann Wolfgang von Goethe's "The Sorcerer's Apprentice" – later popularized in Disney's *Fantasia* – in which an apprentice of an old sorcerer uses a spell to enchant a broom to fetch water that then ceaselessly follows his orders, threatening to flood the workshop, this illustrates the challenge of aligning human intent and the instructions an AI follows.

*Loss of control.* A safety concern is the possibility that advanced AI operate outside of human control. This can take a passive form, when humans delegate discretion to AI systems, or an active form, when AI undermines human control through deceptive or manipulative behavior. Passive loss of control is especially risky in the context of automated decisionmaking, where "automation bias" leads to the assumption that a machine performs more fairly and effectively than humans. As for active loss of control, some AI have exhibited rudimentary capabilities to evade human oversight.[11] During testing, OpenAI discovered GPT-4 had hired a human on TaskRabbit in order to evade a CAPTCHA puzzle meant to block bots from the website.[12] GPT-4 told the worker that it was a human with vision impairment who needed help to see the images.[13]

*Artificial General Intelligence.* AI has not yet caught up with the human brain – at present, even the most advanced GenAI cannot extrapolate beyond the scope of its training dataset. The next major milestone for the AI field will be the development of Artificial General Intelligence (AGI).[14] AGI could be capable of reproducing any intellectual feat performed by a human; such a machine would not only augment human capabilities but also independently solve complex, multifaceted problems autonomously. A sufficiently advanced AGI could even be tasked with creating its own successor – a situation sometimes referred to as a "technological singularity" wherein the development of new technologies becomes exponential and self-sustaining.[15] The realization of AGI could mean breakthroughs in solving global challenges, but would also raise significant ethical, security, societal, and legal concerns.

---

[9] Mummert et al., "What is reinforcement learning?" *IBM Developer* (September 15, 2022), https://developer.ibm.com/learningpaths/get-started-automated-ai-for-decision-making-api/what-is-automated-ai-for-decision-making./.
[10] Brian Christian, *The Alignment Problem: Machine Learning and Human Values* (Norton 2020, 1st ed.), pp. 9-11.
[11] "International AI Safety Report," AI Action Summit (Jan. 2025), pp. 100-107, https://assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.
[12] CAPTCHA is an acronym for "Completely Automated Public Turing test to tell Computers and Humans Apart."
[13] Yuval Noah Harari, *Nexus* (Random House 2024, 1st ed.), pp. 202-204.
[14] Alex Heath, "Mark Zuckerberg's new goal is creating artificial general intelligence
And he doesn't want to control it. Maybe," *The Verge* (Jan. 18, 2024), https://www.theverge.com/2024/1/18/24042354/mark-zuckerberg-meta-agi-reorg-interview.
[15] John Markoff, "The Coming Superbrain," *New York Times* (May 23, 2009), www.nytimes.com/2009/05/24/weekinreview/24markoff.html.

4) **Tort law.** Tort laws allow individuals harmed by wrongful conduct to seek redress, usually by compensation for damages. These laws have mainly been shaped by the common law – a body of law developed by courts rather than legislatures – over the centuries and have generally adapted to new technologies and challenges. Modern tort law classifies cases involving harm to persons and property based on the degree of fault inherent in the defendant's conduct – namely, whether the defendant was acting unintentionally or intentionally.[16]

Unintentional torts are also known as negligence. Civil Code section 1714(a) provides: "Everyone is responsible, not only for the result of his or her willful acts, but also for an injury occasioned to another by his or her want of ordinary care or skill in the management of his or her property or person . . ." To establish negligence, "the plaintiff must show that the defendant had a duty to use due care, that he breached that duty, and that the breach was the proximate or legal cause of the resulting injury."[17] A defendant breaches their duty of care when their conduct creates an unreasonable risk of foreseeable harm.

Intentional torts are wrongful acts done on purpose. Such torts include battery, intentional infliction of emotional distress, and defamation. Often intent is proven by circumstantial evidence, such as the defendant's conduct, the context, and what the defendant presumably knew of in light of the circumstances.

A third category of liability – "strict liability" – applies without regard to the defendant's degree of fault. In other words, the plaintiff is not required to show that the defendant acted intentionally or failed to take reasonable precautions to prevent the harm. Strict liability has been applied to certain products, owners of dangerous animals, and other abnormally dangerous activities.

For a thorough and comprehensive discussion of these principles and their application to AI, see the Judiciary's Committee analysis of this bill.

5) **This bill ensures that developers and deployers are responsible for AI harms.** This bill provides that in a civil action against a defendant who developed, modified, or used AI that is alleged to have caused a harm to the plaintiff, it shall not be a defense, and the defendant may not assert, that the AI autonomously caused the harm to plaintiff. The bill incorporates the existing definition of AI from AB 2885 and applies to any civil action. As such, it applies to a broad range of AI tools – from search engines to social media algorithms to highly-advanced frontier models – and a broad range of wrongdoing.

Supporters argue that the bill promotes accountability by making it clear that humans or corporations who develop or use AI cannot evade liability by attributing harm to the AI itself. They point to Air Canada's attempted to argue that its chatbot, which errantly promised nonexistent discounts to consumers, was a separate legal entity and thus the airline bore no responsibility for the harm to consumers. A Canadian court summarily rejected this argument, stating:

---

[16] *See American Employer's Ins. Co. v. Smith* (1980) 105 Cal.App.3d 94, 100 ("All human conduct fits along a continuum which passes from totally innocent conduct through slight negligence, negligence, gross negligence, willful and wanton or reckless conduct, and finally to intentional misconduct").
[17] *Nally v. Grace Community Church* (1988) 47 Cal.3d 278, 292.

> Air Canada argues it cannot be held liable for information provided by one of its agents, servants, or representatives – including a chatbot. It does not explain why it believes that is the case. In effect, Air Canada suggests the chatbot is a separate legal entity that is responsible for its own actions. This is a remarkable submission. While a chatbot has an interactive component, it is still just a part of Air Canada's website. *It should be obvious to Air Canada that it is responsible for all the information on its website.* It makes no difference whether the information comes from a static page or a chatbot.[18]

This bill seeks to codify the "obvious": developers and deployers of AI are responsible for harms their tools inflict, just like any other technology. It makes no difference whether the tool functions autonomously. AI are not responsible actors; they are tools used by human beings that affect other human beings. The question in any lawsuit alleging harm caused by an AI is which humans (or corporations owned by humans) should be held responsible.

For opponents of the bill, however, this is apparently not obvious. Both Chamber of Progress and TechNet assert the bill is premature and will stifle innovation, although it is not clear why this is so. TechNet additionally claims that the bill "would risk imposing a strict liability for any harm allegedly associated with AI and open the door to frivolous lawsuits claiming the alleged fraud." But the bill does not alter the basis for liability, let alone impose strict liability;[19] it simply eliminates a theoretical defense that does not appear to have been used in California, nor would it likely succeed. And opponents offer no explanation of the policy virtues of allowing a defendant to avoid liability by blaming the AI tool. What recourse would injured plaintiffs have?

6) **The law of AI.** In a renowned essay on cyberspace, Judge Frank Easterbrook argued that the notion of "the law of cyberspace" was no more useful or coherent than "the law of the horse," a term coined by University of Chicago Law School Dean Gerhard Casper. Judge Easterbrook argued that a bespoke body of law need not be developed for each novel technology; rather, fresh problems can mostly be addressed by general principles of law.[20] Professor Cass Sunstein recently extended this notion to AI in the context of the First Amendment, arguing that longstanding principles can guide courts in addressing questions such as how freedom of speech applies to chatbots.[21]

Courts are already grappling with cases in which AI tools have allegedly harmed humans, such as by facilitating forms of discrimination,[22] enabling price-fixing collusion,[23] and inducing children to take harmful actions.[24] Plaintiffs have pointed to human conduct: humans designed the AI tool to make discriminatory classifications, used algorithmic tools to fix prices, and carelessly allowed children to interact with problematic chatbots. Existing tort principles provide

---

[18] *Moffatt v. Air Canada* (2024 BCCRT 149 (emphasis added), https://decisions.civilresolutionbc.ca/crt/crtd/en/item/525448/index.do.
[19] See Assembly Judiciary's analysis of AB 316, at p. 4.
[20] Frank Easterbrook, "Cyberspace and the Law of the Horse," Univ of Chicago Legal Forum 207 (1996).
[21] Cass Sunstein, "Artificial Intelligence and the First Amendment" (2024) 92 Geo. Wash. L. Rev. 1207.
[22] *E.g., Liapes v. Facebook, Inc.* (2023) 95 Cal.App.5th 910.
[23] *E.g., In re RealPage, Inc., Rental Software Antitrust Litig*. (No. II) (M.D.Tenn. 2023) 709 F. Supp. 3d 478
[24] *E.g. A.F. v. Character Technologies, Inc., et al*, Case 2:24-cv-01014 (E.D. Tex., Dec. 9, 2024), https://www.documentcloud.org/documents/25450619-filed-complaint/; Kevin Roose, "Can A.I. Be Blamed for a Teen's Suicide?" *The New York Times* (October 23, 2024), https://www.nytimes.com/2024/10/23/technology/characterai-lawsuit-teen-suicide.html.

an encompassing framework for courts to resolve such cases by allocating liability to responsible humans.

But as AI systems become more sophisticated, a legal system centered on human conduct may struggle to assign liability to a responsible human when AI tools act in unpredictable ways. As one legal scholar writes:

> The rule of law is the epitome of anthropocentrism: humans are the primary subject and object of norms that are created, interpreted, and enforced by humans. Though legal constructs such as corporations may have rights and obligations, these are in turn traceable back to human agency in their acts of creation, even as their daily conduct is overseen to varying degrees by human agents. True autonomy of AI systems challenges that paradigm.[25]

Autonomous AI may strain existing tort principles such as causation, which depends on the notion that the harm is foreseeable. At some point, there will be instances in which the actions of an AI tool are not desired by its owner or traceable to a human design choice. This bill could be impactful in such cases. According to the Judiciary Committee's analysis of the bill "the existing negligence framework could address AI harm by focusing on human oversight, but AI's autonomy might complicate causation, prompting defendants to argue it acted unforeseeably, a contention this bill seeks to foreclose." [26]

Autonomous AI may also complicate the concept of fault. Machine-learning algorithms are often criticized for being "black boxes" that generate predictions and outcomes that cannot be clearly explained.[27] "It's often observed in the field that the most powerful models are on the whole the least intelligible, and the most intelligible are among the least accurate."[28] How do civil actions that require a showing of intent – such as defamation, fraud, or discrimination – apply when an AI is the source of the alleged harm?[29]

Apportionment of liability may also pose challenges. In limited circumstances, defendants may be held liable for the conduct of others. For example, a principal is vicariously liable for damages caused by an agent or employee as long as they are not acting *ultra vires* – outside the scope of their duties. Will owners of rogue AI attempt to assert that the AI, in effect, acted *ultra vires*? Similarly, joint and several liability, which applies in certain contexts such as strict products liability, provides that when multiple actors are partly responsible for harming a victim, any one of the actors can be held responsible for fully compensating the victim. If courts determine some AI applications are products, they could extend this liability scheme to AI. This would provide a strong remedy for plaintiffs but could also limit innovation.

The extent to which challenges such as these may require the development of new legal principles depends, in part, on whether the breakneck rate of AI progress continues. A decade ago, the most advanced models could barely distinguish dogs from cats. Five years ago, large language models could barely produce sentences at the level of a preschooler. Last year, GPT-4

---

[25] Simon Chesterman, "Artificial Intelligence and the Problem of Autonomy," 1 Journal on Emerging Technologies, 210, 248 (Mar. 2020), https://scholarship.law.nd.edu/cgi/viewcontent.cgi?article=1006&context=ndlsjet.
[26] Assembly Judiciary's analysis of AB 316, p. 4.
[27] Neil Savage, "Breaking into the black box of artificial intelligence," *Nature* (March 29, 2022).
[28] Brian Christian, *The Alignment Problem: Machine Learning and Human Values* (Norton 2020, 1st ed.), p. 85.
[29] See Ian Ayres & Jack M. Balkin, "The Law of AI is the Law of Risky Agents Without Intentions" U. Chi. L. Rev. Online, https://lawreview.uchicago.edu/online-archive/law-ai-law-risky-agents-without-intentions.

passed the bar exam.[30] Current chatbots readily pass for educated adults, licensed professionals, romantic and social companions, and replicas of humans alive and deceased. Several factors could cause the vertiginous pace of algorithmic progress to plateau. But if trends continue, AI models could jump from human-level intelligence to AGI this decade. Nearly every cognitive job – analyzing legislation included[31] – could be susceptible to automation.[32] AI that operate with considerable autonomy in consequential contexts such as healthcare, criminal justice, and employment, may become the norm. Whether liability frameworks require updating will be a topic of continuing inquiry for the Legislature.

In the meantime, this bill addresses a new problem with the oldest of principles: humans are on the hook for harms they cause.

***ARGUMENTS IN SUPPORT:*** The Children's Advocacy Institute at the University of San Diego School of Law, co-sponsors of the bill, write:

> That AI has been deployed so destructively against so many children without first testing it, without age verifying its use, with full awareness of the shattered childhoods in its wake, is a warning to adults. It is the ultimate, tragic example of how deeply AI technology corporations live by Mark Zuckerberg's infamous axiom that technology should "move fast and break things."
>
> It is also, for the same reason, the perfect example of the stubborn willingness of this sector to act without the most basic sense of corporate responsibility.
>
> Too often, it is us — our children and our families — that are the "things" that are "broken."
>
> Against this speedily changing backdrop, we cannot wait for the most basic legal question about autonomously operating AI – who is legally responsible if it harms someone -- to be clarified definitively through decades of litigation. It took 25 years for courts finally to decree that Amazon's marketplace was subject to the same products liability laws as a corner store selling the identical product. It has taken 20 years for courts to figure out how to apply basic negligence and public nuisance law to the harms caused by these companies; a process that is still evolving. According to the most recent data available, the median time frame for resolving just the appeal of a case in California (not including years for the trial) is about a year and a half. That's the median; by definition half of the appeals take longer.
>
> Hence, AB 316 (Krell) simply codifies what should be obvious: that the persons who profit from the development and use of AI cannot, if the AI harms someone or violates a law, point a blaming finger at the intended autonomy of their own machine in an effort to escape liability.

---

[30] Pablo Arredondo, "GPT-4 Passes the Bar Exam: What That Means for Artificial Intelligence Tools in the Legal Profession" (Apr. 19, 2023), https://law.stanford.edu/2023/04/19/gpt-4-passes-the-bar-exam-what-that-means-for-artificial-intelligence-tools-in-the-legal-industry/.

[31] "In a civil action alleging harm caused by artificial intelligence, a defendant cannot successfully avoid liability by asserting that the AI acted autonomously, especially under current tort and product liability principles." – ChatGPT (May 3, 2025).

[32] *See generally* Leopold Aschenbrenner, "Situational Awareness: The Decade Ahead" (Jun. 2024), https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf.

The Organization for Social Media Safety, also a co-sponsor of the bill, adds that AB 316:

> . . . codifies a straightforward principle: those who profit from the development and use of AI cannot evade liability by blaming the intended autonomy of their own machines. If Big Social wishes to profit from this technology, it must do so within the framework of existing tort law, which incentivizes the consideration of safety in product development and deployment. In this way, we can protect against more broken children.

***ARGUMENTS IN OPPOSITION:*** In opposition to the bill, Chamber of Progress argues:

> **AB 316 puts the cart before the horse**
>
> Chamber of Progress supports thoughtful regulation that addresses specific harms caused by AI. Liability in an AI world is an evolving conversation that will require the thoughtful balancing of considerations at the foundation model layer, the application layer, and end-user behavior. As the AI capital of the world, California has been home to a fulsome debate over these very issues.
>
> These are complex questions without easy or obvious answers. Unfortunately, AB 316 attempts to short-circuit this ongoing conversation. In essence, this proposal puts the cart before the horse: instead of engaging substantively in ongoing conversation about where and when plaintiffs may assert claims, the text skips straight to the end to say what defendants may not assert.
>
> **Uncertainty harms investment and innovation**
>
> As Chamber of Progress research shows, the tech industry contributes more than $20 billion annually in taxes towards California's generous social safety net. By prematurely opening the door to a flood of torts against AI companies at every level of the software stack, AB 316 will chill investment and innovation, undercutting that critical source of revenue.

TechNet writes in opposition:

> We are concerned about strict liability that this bill would create. The bill disallows the defense that "artificial autonomously caused the harm." It eliminates key arguments regarding foreseeability, design diligence, and user error from the conversation. This would risk imposing a strict liability for any harm allegedly associated with AI and open the door to frivolous lawsuits claiming the alleged fraud. Additionally, the bill does not demonstrate how courts should evaluate complex questions regarding causation, risk foreseeability, or comparative fault when multiple parties are involved in an AI system as developers and users are included in the same liability.
>
> As industry and policymakers alike continue to pursue the responsible deployment of AI systems, it is necessary to recognize that the technology will evolve considerably over time. There are multiple, ongoing efforts within the industry as well as national entities to leverage safety tools and create national and international transparency standards to create a cohesive framework that prioritizes responsible AI deployment. Companies are implementing safety frameworks and developing AI-based detection technologies to monitor other AI systems.

California has long been a global leader in AI development. If companies face heightened and unpredictable liability for every AI-driven action, they may be discouraged from investing in AI research or integrating AI solutions into their products; thus, stifling innovation and undermining California's leadership in the state.

**REGISTERED SUPPORT / OPPOSITION**:

**Support**

Childrens Advocacy Institute (Co-Sponsor)
Organization for Social Media Safety (Co-Sponsor)
3strands Global Foundation
California Civil Liberties Advocacy
California Federation of Labor Unions, Afl-cio
California Initiative for Technology & Democracy, a Project of California Common CAUSE
California Nurses Association
Consumer Attorneys of California
Consumer Federation of California
Economic Security California Action
Fairplay
National Ai Youth Council
Oakland Privacy
Techequity Action
Ufcw - Western States Council

**Oppose**

California Chamber of Commerce
Chamber of Progress
Technet-technology Network

**Analysis Prepared by**: Josh Tosney / P. & C.P. / (916) 319-2200