




GFELLER  LAURIE^{LLP}
ATTORNEYS AT LAW

Artificial Intelligence, Legal Liability, and Insurance

Vincent J. Vitkowsky



Connecticut Massachusetts New York New Jersey Pennsylvania

Artificial Intelligence, Legal Liability, and Insurance

Executive Summary

In the last year, Artificial Intelligence captured the world’s attention. This followed the public release of ChatGPT, a specific kind of Artificial Intelligence (“AI”) known as Generative AI, and more specifically a “large language model-based Chatbot” capable of generating human-like text and having conversation-like interactions with humans. Chatbots are an extension of a broader technology that has existed for many years and has been used in many applications, in many areas, including healthcare, finance, and engineering. Almost every industry and human endeavor can use some form of AI, and the scope of potential applications has expanded significantly with recent advances in Generative AI.

AI systems vary widely, and the more autonomous they become, the more unpredictable they become. In many complex systems, it is not possible to trace the process, logic, and analysis giving rise to a particular output. This is referred to as Black Box AI. For a legal system heavily reliant on determining causation, this poses problems in assigning liability.

Large language-model Chatbots pose other concerns. Simply put, they scan massive amounts of textual content, review them, assign probabilities to likely word sequences, and predict the most likely text in response to a question. They are designed to provide answers, and they do that, even if it means providing answers that are not true. This is referred to as “hallucinations” or “stochastic parroting.” Moreover, many are not capable of reasoning. For example, they cannot solve basic logic puzzles. And they entirely lack judgment, moral or otherwise.

The use, misuse, or malfunction of AI systems can give rise to claims for all manner of financial loss, personal injury, property damage, or legal infractions. Potentially liable parties include everyone in the chain of creation, production, and operation of AI systems, and also the end users.

The legal rules that will apply to AI are in their early development. Sources of liability will be drawn from contract law and tort law, and often professional liability and product liability law. The special factors implicated in AI systems will require adaptation and refinement of established principles, and the development of new duties and standards of care. Matters such as the presentation of technical evidence, the allocation of liability along the supply chain, manufacturing, and approval process, and the proof of causation will all be much more complex. The credibility of evidence may be subject to challenge on the grounds it is false evidence created by AI systems.

Potential insurance coverage for AI-related risks and liabilities may be found in certain traditional insurance policies, perhaps as “Silent AI” in the absence of specific exclusions. These principally include professional liability, tech e&o, miscellaneous liability, commercial general liability and products hazards, property, and cyber insurance policies.

However, there are many clear gaps in each of these policies. To meet coverage needs, a few companies are offering new specialist insurance products which offer some combination of coverage found in various traditional policies, as well as new coverages.

All insurers should examine the impact of AI in a comprehensive review of their portfolios.

Vince Vitkowsky
Gfeller Laurie LLP
November 1, 2023
www.gllawgroup.com

Please note that this Paper is for informational purposes only and is not comprehensive. It does not constitute the rendering of legal advice or opinions on any subject matter. The distribution of this Paper to any person does not constitute the establishment of an attorney-client relationship.

Artificial Intelligence, Legal Liability, and Insurance

Understanding Artificial Intelligence

The release of ChatGPT by the OpenAI research laboratory in November 2022 has brought widespread interest in “Artificial Intelligence” (“AI”). But ChatGPT is only one application of a specific type of AI, known as “Generative AI.” More specifically, ChatGPT is a “large language model-based Chatbot,” capable of generating human-like text and having conversation-like interactions with humans. It came into broad use upon its release as a free software application available to the public. It is now available in two versions, ChatGPT 3.5, which remains free, and a further version, ChatGPT 4.0, which is available to paid subscribers. Several leading technology companies have released their own competing versions.

These Chatbots are an extension of a broader technology, and related technologies, which have existed for many years. Several terms have been used to describe the technologies. The terms have been used inconsistently and imprecisely, and often overlap. The FTC said in a report to Congress in 2022 that “AI is defined in many ways and often in broad terms.¹ Still, it is helpful to provide some definitions, as a foundation for understanding the matters addressed in this Paper. One generally accepted source of definitions is the online dictionary Technopedia.

Artificial Intelligence. Techopedia defines Artificial Intelligence as follows:

Artificial intelligence (AI), also known as machine intelligence, is a branch of computer science that focuses on building and managing technology that can learn to autonomously make decisions and carry out actions on behalf of human beings AI is not a single technology. Instead, it is an umbrella term that includes any kind of software or hardware component that supports machine learning ..., expert systems, generative AI ... and certain types of robotics.²

“Expert systems” are of special interest. These were developed in the 1970s and 1980s, and “aimed to capture the knowledge and decision-making processes of human experts in specific domains and use that knowledge to provide recommendations or make decisions.” They “have practical applications in healthcare, finance, and engineering.”³

Machine Learning (ML). Technopedia defines Machine Learning (herein, “ML”) as “a subset of AI that focuses on building models ... that allow computers to learn from data. Instead of being specifically programmed to perform a task, ML models use data to make predictions or decisions.”⁴

¹ F.T.C., COMBATTING ONLINE HARMS THROUGH INNOVATION: FEDERAL TRADE COMMISSION REPORT TO CONGRESS 1 (2022).

² Margaret Rouse, *Artificial Intelligence*, TECHNOPSIS (last updated Aug. 15, 2023), <https://www.techopedia.com/definition/190/artificial-intelligence-ai>.

³ *Id.*

⁴ Margaret Rouse, *Machine Learning (ML)*, TECHNOPSIS (last updated Aug. 9, 2023), <https://www.techopedia.com/definition/8181/machine-learning-ml>.

Technopedia elaborates by explaining that ML “builds algorithmic models to identify patterns and relationships in data. In this context, the word *machine* is a synonym for computer program and the word *learning* describes how ML algorithms ... become more accurate as they receive additional data.”⁵

Generative AI. Technopedia defines Generative AI as “a broad label that’s used to describe any type of artificial intelligence ... that can be used to create new text, images, video, audio, code, or synthetic data.”⁶

Training. Training is a crucial concept in all AI systems. Training refers to the data sets that an AI system includes and uses to perform its functions. The practical application of AI systems in business and many other areas was empowered by the emergence of the internet and advances in big data analytics --- “the digitalization of everything.” The data sets come from public databases and data licensed from third-party providers. In most cases, the most valuable component of an AI system is not the software. The main value is in the proprietary training data set.

Uses and Applications. Some AI systems have been in use for many years, even decades. AI systems of varying complexity have become ubiquitous. To cite a few examples, they are used in GPS systems, autonomous vehicles, drones, workplace robots, robotic security guards, medical imaging and diagnostics, financial modelling, portfolio management, virtual assistants, image and facial recognition, personalized shopping and other recommendation systems, inventory management, customer services, and back-office and compliance functions. Financial institutions use regulatory technology (“regtech”) to assist in compliance functions such as onboarding, anti-money laundering and fraud detection. Generative AI systems are also used in creating brochures and other marketing deliverables, customer experience management, employment decisions, and web publishing.

The development of Generative AI has expanded the scope of potential use in ways that can have broad and sometimes critical impacts on many aspects of life. For example, Generative AI can create fictional or non-fictional text, digital images, video and audio content. It can be used in portfolio management to make predictions based on past performance and current trends. It has new applications in healthcare, in everything from diagnosing conditions to developing personal treatment plans.⁷ It can create amazingly believable deepfakes. It can be used to facilitate cyberattacks and cybercrime.

Considerations Concerning AI Systems Generally

AI systems vary widely. Their capabilities depend on what they are designed to do, and the nature and quality of their training data set. Some systems are limited to their original training data set. Others can acquire new data which is added automatically to their training data set

⁵*Id.*

⁶ Margaret Rouse, *Generative AI*, TECHNOPEdia (last updated Jun. 27, 2023), <https://www.techopedia.com/definition/34633/generative-ai>.

⁷ *Id.*

when it appears on the internet. Others can have their data set increased by the output of other AI systems, as well as their own additions or conclusions.

AI systems can be automated” or “autonomous.” “Automated AI” is pre-programmed and designed to perform specific tasks, in which a human retains overall control of its operation. “Autonomous AI” performs its functions without human intervention, input, or direct supervision, and is able to predict, plan, and be aware of the surrounding environment.

Autonomous AI systems can produce some surprising results. For example, there is a report of a collegiate robotics competition in which the robots were programmed to compete to gather the most “sheep” into their individual pens. The robots had to develop their own tactics. One robot threw a single sheep into its pen and shut the gate. It then set upon destroying the other robots, so no other sheep could be captured.⁸ One could say that even though the robot had not been programmed for deviousness, it acquired that trait on its own.

There is another critical consideration. “Sophisticated AI methods produce results without explaining why or how their process works.”⁹ Often, it is not even always *possible* to trace the process, logic, and analysis gave rise to a particular output, or what parts of the training data set were used to generate a particular output. Sometimes a data scientist or ML engineer is able to explain how a specific AI output was arrived at (for example, when the algorithm is coded to log its decision-making process). In those cases, the AI/ML model is referred to as “Explainable AI.” But that is not always possible, and in those instances, the AI/ML model is referred to as “Black Box AI.”¹⁰ Most advanced AI/ML models have significant Black Box aspects. For a legal system heavily reliant on determining causation, this poses problems in assigning liability.

Special Considerations Concerning Generative AI Chatbots

ChatGPT and its competing large language-model Chatbots can be used to generate text. They respond to a question, or “prompt,” from a user. They have been trained on a massive amount of textual content contained in the digital world. They review these at otherworldly speeds, assign probabilities to likely word sequences, and predict the most likely text.

The key point to understand about the current generation of AI Chatbots is that they are not especially intelligent. Or more precisely, they have a limited kind of intelligence. They can possess enormous knowledge, potentially including everything ever written that is found in the digital world. But they cannot evaluate the extent to which what is written is true or false. Put differently, they possess knowledge, but not wisdom. They cannot exercise judgment. Simply, they are designed to provide answers to prompts. They do that, even if it requires making up facts. The result is that the answers may not be true. That is referred to as “hallucinations” or

⁸ Sahara Shrestha, *Nature, Nurture, or Neither?: Liability for Automated and Autonomous Artificial Intelligence Torts Based on Human Design and Influences*, 29 Geo. Mason L. Rev 375, 397 (2021) (citing Luba Belokon, *Creepiest Stories in Artificial Intelligence (AI) Development*, MEDIUM (Sept. 21, 2017), <https://perma.cc/UD8T-FBFR>).

⁹ Henry Kissinger et al., *ChatGPT Heralds an Intellectual Revolution*, WALL ST. J. (February 24, 2023), <https://www.henrykissinger.com/articles/chatgpt-heralds-an-intellectual-revolution/>.

¹⁰ Rouse, *supra* note 4.

“stochastic parroting.” “What triggers these errors and how to control them remains to be discovered.”¹¹

There is the well-known example of a law firm which was sanctioned for filing a brief generated by ChatGPT which included completely fabricated case names and citations, with fictitious holdings and reasoning.¹²

In another example, experimenters asked ChatGPT to provide “six references on Henry Kissinger’s thoughts on technology.” It generated a list of articles with plausible topics and appearing in plausible publications. Only one was real (but with an incorrect date). The others were “convincing fabrications.”¹³

A second key point is that Chatbots are not capable of reasoning. A Wall Street Journal opinion article provided an example:

Still, AI has a long way to go. A recent study typed this prompt into dozens of large language models: ‘Sally (a girl) has 3 brothers. Each brother has 2 sisters. How many sisters does Sally have?’ For those who like these mental puzzles, the correct answer is one. None got it. ChatGPT 4.0 got closest with ‘Sally has 2 sisters.’ Others answered three or six. Generative AI is still far from civilization risk.¹⁴

This example is so striking that this Paper’s author consulted with a leading AI scientist to ask whether it could possibly be true. The response was “absolutely yes.” These systems cannot reason. They are not wise. They are designed to scan their training data set, search for words and word orders, predict the most commonly used order, and then recite them in an answer. Again, how they choose from among the potentially infinite universe of words to generate their precise response is not known.

Finally, Chatbots and other Generative AI systems lack judgment. Unlike humans (or at least *some* humans), they are not moral, empathetic, strategic, or visionary entities.¹⁵

¹¹ Kissinger, *supra* note 9.

¹² Sara Merken, *New York lawyers sanctioned for using fake ChatGPT cases in legal brief*, REUTERS (June 26, 2023, 4:28 AM), <https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>.

¹³ Kissinger, *supra* note 9.

¹⁴ Andy Kessler, *Schumer Wants a Cut of AI*, WALL ST. J. (September 18, 2023), https://www.andykessler.com/andy_kessler/2023/09/wsj-schumer-wants-a-cut-of-ai.html.

¹⁵ In reading the previous sections, readers should understand that I am not a technologist, possess no coding or data analysis ability, and have average mathematical ability. My skill set is that of a practicing lawyer: analytical reasoning and expository narration. I have tried to apply those to the very technical matters discussed above for the benefit of readers with similar skill sets who need to grapple with the legal issues raised by complex AI systems.

The Legal Framework

The legal rules that will apply to AI are in their early development. Basic principles can be drawn from existing contract and tort law, especially professional liability and product liability law. But many cases will present new variations and considerations, requiring novel analytical applications and refinements. What some of the approaches might be, or should be, are the subject of current scholarly analysis and debate.

This is especially true with Black Box AI systems (see discussion above), which include many or most advanced AI systems. This means neither the creators nor users can fully understand the process, justification, and analysis leading to the output. Some have argued that “when an AI entity inflicts harm or injury, the black box, then, whose decisions are ultimately inexplicable, posing a problem for assigning liability.”¹⁶ A full evaluation and discussion of this highly technical subject is beyond the scope of this Paper. For present purposes, it is sufficient to note that this consideration will complicate the process of allocating liability.

Potential Claims and Liable Parties

Claims. The use, misuse or malfunction of AI systems can cause all manner of financial loss, personal injury, property damage, or legal liability. For example, runaway robots injure people and property, GPS systems and autonomous vehicles cause crashes, and medical diagnostic devices provide inaccurate outputs.

They can also lead to regulatory claims. For example, the Australian Bank Westpac was fined Aus\$1.3 billion by its regulators, largely related to AI system programming errors which led to 23 million breaches of financial crime laws over five years.¹⁷

Generative AI large language-models Chatbots and other forms of Generative AI present an expanded set of potential claims. Principally, these include intellectual property infringement, employment related claims such as inherent bias in hiring (based on the training data set), breaches of the user’s own confidential information or employee data, breaches of cybersecurity, violations of comprehensive state data privacy laws, and publication of slanderous and defamatory materials. Moreover, they can be used to develop amazingly believable deepfakes, malware, ransomware, phishing attacks, and other tools that facilitate cybercrime.

Potentially Liable Parties. Liability for losses may be imposed on, among others, (1) any entities involved in the creation, software, design, engineering, installation, and maintenance of an AI system, (2) the creator of the training data sets, who may be different, (3) manufacturers, (4) AI consultants, and (5) the entity which operates the AI system (collectively, “AI System Providers”). Liability may also be imposed on the end users.

¹⁶Anat Lior, *Artificial Intelligence and Tort Law – Who should be Held Liable when AI Causes Damages?*, HEINRICH BÖLL STIFTUNG (December 24, 2021), <https://il.boell.org/en/2021/12/24/artificial-intelligence-ai-tort-law-and-network-theory-who-should-be-held-liable-when-ai>.

¹⁷Ross P. Buckley et al., *Regulating Artificial Intelligence in Finance: Putting Humans in the Loop*, 43 Sydney L.J.43, 54 (2021) (also recommending a series of internal governance tools to minimize potential liability from AI-related issues).

Sources of Liability

In the first instance, AI System Providers all along the chain may seek to limit or allocate liability by contractual provisions or disclaimers of liability. Given the novelty of the technology and the causation issues, it is far from clear how such attempts will fare when tested in the courts.

Apart from contract language, courts will need to grapple with their state's Uniform Commercial Code ("UCC"), which governs contracts for the sale of goods.

Contract Law

A particular challenge here is that courts diverge on whether software is a "good" or a "service." Software is frequently but not always found to be a good, because generic software comes in tangible form. Where the UCC applies, it creates certain warranties.

Express warranties, which are breached when the defendant makes a false, material, factual statement about the product that the plaintiff relied on.

Implied warranties of merchantability, which are breached when the product was unsuitable for the ordinary purpose for which it is used.

Implied warranties of fitness for a particular purpose, which are breached when the product was required for a special purpose, the plaintiff relied on the defendant's skill or judgment in selecting the product for that purpose, and the product could not serve that purpose.

Beyond the seller and buyer, in some instances these warranties extend to injury to the person or property of third parties.

It is also possible that a complex AI system might not be deemed to be a generic good. In fact, many AI systems are likely to be customized for specific uses -- which could often be characterized as "services" -- by particular customers. Here, the "Predominant Factor" test would come into play. Courts evaluate whether the predominant purpose of the transaction was the development of the software, or the services provided. There will be many instances in which the answer is not readily apparent. Courts will look to the degree of customization for the specific purpose in determining whether the AI system was a good or a service, and whether the warranties apply. If the AI system is deemed to be a service, negligence rules apply. The rules will be developed over time, in courts throughout the country, evaluating many different AI systems and sales.

Tort Law

Many AI-related claims will be resolved by applying tort law, as it develops. The most basic tort principle is negligence. It imposes liability if a defendant has failed to meet the standard of care that a reasonable person should exercise in the circumstances. It can be applied to both AI System Providers and end users. Beyond basic negligence, the main strains of tort law that will be implicated are professional liability and product liability.

For the sake of illustration, this Paper will review the potential liability of medical providers. Hundreds of AI-driven medical devices have been approved by regulators. They principally involve imaging and diagnostics, although some advanced AI systems may also develop treatment plans. There is potential liability by healthcare providers, including human beings such as physicians, nurses, or systems technicians. Their liability will be measured by professional liability and negligence standards. The core issue often will be the applicable standard of care. Questions would include whether the professionals performed independent critical analyses of the AI output, and whether they seemed reasonable based on their experience. But the essential point is that no standard of patient care when using AI medical devices has yet been established.

Hospitals and other healthcare facilities may be liable through vicarious liability under the *respondeat superior* doctrine. In addition, there will be questions about the process used to purchase and evaluate specific AI systems, including whether the hospital or facility has an independent duty to evaluate or seek advice on the quality of the technology and algorithms. Again, at present, there are no established standards of care for hospitals and facilities, so each new case will require a new analysis.

In litigation arising from all kinds of AI systems, another group of potentially liable entities are the AI System Providers. Their exposure is subject to the law of Product Liability. On one level, this will simply require applying accepted principles. But on another, there will need to be considerable further development of these principles because of the special factors implicated by AI systems. The presentation of technical evidence, the allocation of liability along the supply chain, manufacturing, and approval process, and the proof of causation will all be much more complex. Indeed, the credibility of evidence may be subject challenge on the grounds that is false evidence created by AI systems.

As an alternative cause of action, in some jurisdictions and in some cases, strict liability might be applied. Although there are some differences in application across states, most states apply Section 402A of the Restatement (Second) of Torts. Under the Section, a defendant who sells a product in a defective condition that is “unreasonably dangerous” may be liable for physical harm or property damage even if it exercised all possible care in the preparation or sale of the product, and even if the injured party had no contractual relationship with the defendant.¹⁸ Whether a system is found to be “unreasonably dangerous” will vary, AI system-by-AI system and case-by-case.

¹⁸ See, e.g., *Taylor v. Intuitive Surgical, Inc.*, 389 P.3d 517, 520 (Wash. 2017) (holding that strict liability was the proper test for a manufacturer’s failure to warn of the potential flaw in its surgical robot).

Potential Insurance Coverage

Traditional Insurance Coverage. Some traditional insurance products may cover some AI-related risks and liabilities. In the absence of specific exclusionary language, to borrow and adapt a phrase from cyber insurance, there could be “Silent AI” coverage. But there will be many gaps.

Professional Liability, Tech E&O, and Miscellaneous Liability policies may afford coverage to medical professionals, lawyers, accountants, financial advisors, consultants, AI System Providers, and a broad range of other professionals.

Commercial General Liability (CGL) policies and their products hazard provisions may afford coverage for bodily injury and property damage to third-parties. CGL Coverage B may provide coverage for invasion of privacy, defamation or trade disparagement, use of others’ advertising ideas, and certain intellectual property infringement in the context of the insured’s own advertising. But it is not entirely clear whether these would apply in the context of AI-related claims, unless expressly addressed in the policies. It is also likely that insureds may try to squeeze a wide range of other claims into CGL policies, which would provide fertile ground for litigation.

Property policies may provide coverage for the insured’s own property damage and business interruption claims.

Cyber insurance policies provide some additional coverage. Generally, they will cover information security and data privacy liability, and business interruption. They might, in theory, cover a case in which a Generative AI output resulted in a breach of *another* entity’s confidential or private information. If they provide reputational damage coverage, as some but not all do, they may cover brand damage when, for example, a smart conversational bot wanders off the reservation or hallucinates wildly. If they are comprehensive standalone cyber policies with a media liability section, many could cover invasion of privacy, defamation, infringement of copyright and trademark, piracy, plagiarism, and misappropriation of ideas occurring on the internet. Some have expanded coverage to include, among other things, false advertising, software copyright infringement, and economic harm to third parties relying on false or erroneous content. Other comprehensive standalone cyber policies may include Tech E&O coverage and even Professional Liability coverage. In cyber insurance policies especially, coverage or exclusions for claims involving AI systems require attention.

However, cyber insurance policies will not cover losses for bodily injury or damage to physical property.

New Insurance Products

Insurers face difficult challenges when writing policies specifically covering AI systems. There is no historical data, and the factual matrix is developing at warp speed.

Yet new insurance products are beginning to emerge. Sophisticated companies rely on their own expertise and proprietary models to set rates. Not surprisingly, the initial products are coming from or are backed by large, long-established multiline insurers and reinsurers who have the resources to gamble on a new line of business. This is not a risk for start-up specialist insurers.

The new policies all differ in specifics, but may provide some combination of coverage for errors and omissions, data breaches, system failures, bodily injury, property damage, and other risks. On behalf of the for-profit National Alliance for Insurance Education & Research, entrepreneur and technologist Amri Tarsis asked ChatGPT to search and list companies offering some form of AI Liability Insurance. As of March 10, 2023, it identified Chubb, AXA XL, Zurich, and Allianz.¹⁹ These identifications have not been verified.

Other large companies have launched specialized products. Since 2018, Munich Re has offered an insurance policy for companies selling AI services. Swiss Re and Chaucer have recently backed Armilla Assurance in what it calls a product warranty, insuring that AI systems perform the way their sellers promised.²⁰

Conclusion

All insurers should examine the impact of AI in a comprehensive review of their portfolios.

Cyber risks and cybersecurity made cyber insurance the “Wild West” of insurance. The rapid development of complex AI systems will lead to risks, liabilities, and products that are in the “Outer Space” of insurance. Fasten your seat belts. This is going to be fast, challenging, and fascinating.

**New York, NY
November 1, 2023**



Vince Vitkowsky is a partner in Gfeller Laurie LLP, resident in New York. He focuses on liabilities, insurance, and litigation of cyber and related risks. Vince assists insurers in all aspects of coverage evaluation and dispute resolution in many lines of business, including cyber, CGL, property, and professional liability. He assists in complex claim evaluations, and at times, the defense of insureds in complex matters. He also assists in portfolio reviews, product development and drafting policies and endorsements. vvitkowsky@gllawgroup.com

Copyright 2023 by Vincent J. Vitkowsky. All rights reserved.

¹⁹ Amri Tarsis, *AI Liability: 5 Key Considerations for Risk & Insurance Professionals*, THE NAT’L ALL. FOR INS. EDUC. & RSCH. (March 10, 2023), <https://scic.com/ai-liability-5-key-considerations-for-risk-insurance-professionals/>.

²⁰ Belle Li, *Is Your AI Model Going Off the Rails? There May Be an Insurance Policy for That*, WALL ST. J. (October 2, 2023), <https://www.wsj.com/articles/is-your-ai-model-going-off-the-rails-there-may-be-an-insurance-policy-for-that-afd012d7>.