

Generative AI and the Law

Halftime Report

BY COLIN E. MORIARTY



This article examines recent developments related to generative AI, including updates on legal challenges, business and policy considerations, and a new ABA rule.

Since they emerged into the public consciousness a few years ago, artificial intelligence models that can produce human-like artwork, music, or text have been the subject of legal battles and public debate. *Colorado Lawyer* published a series of articles in 2023 detailing the early challenges to the legality of generative AI, how generative AI was trained, the business use cases and risks, and legal ethics issued of concern to lawyers using the technology.¹ This article provides updates on how these issues have played out and how they continue to evolve.

Evolving Claims and Decisions in Lawsuits Against Generative AI

Over the last year, legal disputes concerning generative AI have evolved and multiplied. The very first cases challenging how generative AI is trained have proceeded through early motions practice, resulting in some orders that shed light on how courts may view the technology. Authors, publishers, and others have filed many new lawsuits challenging the fundamental legality of the way generative AI is built, with sophisticated complaints from the Recording Industry Association of America and *The New York Times* joining the early class action plaintiffs.² Most of them attack large language models (LLMs).³ While these lawsuits usually assert copyright and copyright-adjacent claims, a few focus on privacy or publicity ideas instead.⁴ At the extreme, some have asserted claims that purport to represent humanity as a whole against the indignity of having their written information scraped to train generative AI or being replaced by robots.⁵

The ultimate fate of generative AI in the legal context is still probably several years away. None of the pending cases appear to have gone to trial or been appealed. The few cases that have resolved appear to have been either stayed in

favor of arbitration or dismissed by stipulation between the parties. In particular, none of them appear to have yet dealt directly with the affirmative defense of fair use, which may be the most important question of all.⁶ Nevertheless, there are some lessons and conclusions to be drawn from the progressing cases.

The Operative Claims Against the Legality of Generative AI Training May Be Limited to Copyright and Contract

The emerging decisions suggest that creative lawyering to assert tort claims or state law claims that overlap with copyright might not have much success. The Northern District of California has held that the Copyright Act preempts other claims like intentional interference with contract, unjust enrichment,⁷ negligence, state law unfair competition,⁸ and the Lanham Act.⁹ This is not surprising given the language of the Copyright Act.¹⁰ It was also the holding from *Thomson Reuters v. Ross Intelligence, Inc.*, a case in the District of Delaware where the court analyzed preemption in the context of Westlaw's Keynotes being scraped to use as fodder for machine learning.¹¹ The Delaware court explained that the relevant inquiry is whether the state law claim has an element that the Copyright Act claim does not, and whether that element makes the claim "qualitatively different."¹² In the case of a claim asserting: "You copied our work, thus interfering with our contracts to license that work," which is similar to the arguments advanced by most plaintiffs against generative AI, preemption applies.¹³

Contract claims in particular, however, seem to be able to dodge preemption. *Thomson Reuters* is an example of this. The court noted that contract claims arising out of violations of licensing or access agreements may not be preempted.¹⁴ *Thomson Reuters*, unlike many of the other anti-AI lawsuits, was not directed

to a model that scraped copyrighted but freely publicly available information off the internet. Instead, it involved a machine learning company, Ross Intelligence, getting access to Westlaw's proprietary "headnotes" and Key Number System, which is only available to paying customers.¹⁵ Ross Intelligence tried to purchase access directly, was turned down, and then hired a third party that already had access to Westlaw and queried the database that way.¹⁶ The Delaware court explained that some of the contract claims were not preempted if they were based on terms that prohibited something different from what was protected by copyright.¹⁷ So, while a clause prohibiting the sale, distribution, or transfer of Westlaw's product was preempted by the Copyright Act, a clause prohibiting access in the form of anti-bot and anti-password sharing was not.¹⁸

Thus, plaintiffs who can claim that their copyrighted material was improperly accessed under the terms of a binding license or contract probably can assert claims other than copyright. Such is the case in the lawsuit against training generative AI on code uploaded to GitHub and made available only under the terms of a license. In *Doe v. GitHub*, the Northern District of California agreed that the plaintiffs stated a claim for breach of contract based on the allegation that scraping and training on the code stored on GitHub violated a license agreement.¹⁹

The availability of independent causes of action not preempted by copyright may be critical. If the copyright claims ultimately fail on fair use or other grounds, these alternative claims may yet survive. Practitioners wishing to protect intellectual property from being used for machine learning should therefore consider devising contracts and licenses that include terms that prohibit or control access or manipulation of material.

Famous Artists May Have Protection Against AI Displacement That Others Lack

Individual artists working in the gig economy who see their style usurped by generative AI may find it difficult to find a remedy, because it may not be possible to copyright a personal style or the sound of a voice.²⁰ But those fortunate enough to already be celebrities and enjoy value associated with their likeness may be able to pursue essentially the same remedy under the theory that the output of generative AI violates their rights of publicity.²¹ Such a claim might not be preempted by the Copyright Law.²² Some commentators suggest that, at least for celebrities or influencers, this is a strong vehicle for protecting creators against replacement by artificial or digital versions.²³

The right to privacy is, however, a state law creation that may vary from state to state. Colorado has arguably not specifically adopted the right to publicity tort in the context of allowing recovery for the commercial exploitation of celebrity,²⁴ though the Colorado Supreme Court has recognized a right against the public appropriation of the name or likeness of another as a form of invasion of privacy.²⁵ The Court has focused more on the idea of protecting privacy than protecting commercial property rights. In *Joe Dickerson & Associates v. Dittmar*, the Court decided that the tort of appropriation in Colorado did not require proof that an individual's name and likeness had preexisting commercial value because mental anguish can be enough to prove damages.²⁶ It stopped short of deciding whether "Colorado permits recovery for commercial damages either under the rubric of privacy or under the right of publicity."²⁷ The Colorado Jury Instructions construes this case as establishing a theory of "public disclosure of private facts"²⁸ and notes only that other jurisdictions allow recovery of damage to commercial use of a persona.²⁹

Some plaintiffs have wielded publicity rights against generative AI with success. An online podcast known as Dudesy created an hour-long audio comedy special purporting to be an AI version of George Carlin.³⁰ Allegedly, one generative AI program, probably an LLM, ingested "five decades of Carlin's original standup comedy routines."³¹ A second generative AI program was

used to learn to pronounce the script using a simulacrum of Carlin's voice and delivery.³² Finally, a third generative AI program was used to create still images of Carlin for advertising the production.³³ The creators did not, however,

“
 Claims based on
 publicity rights
 may avoid some of
 the more complex
 or novel questions
 about generative
 AI because they
 are agnostic as
 to the particular
 tool being used to
 infringe on privacy
 or publicity.”

fraudulently claim the recording to be a legitimate product of Carlin. They were very open that the work was the result of generative AI.³⁴

In addition to the now-standard copyright claims, the entity holding Carlin's intellectual property rights asserted that Dudesy was violating entities' rights of publicity in the

decident.³⁵ Citing California law, the complaint argued that Dudesy's project diluted Carlin's own legacy, could cause customer confusion, and represented an improper exploitation of Carlin's image, personality, and voice for profit.

The court, however, did not reach the merits of the claim for violation of the right of publicity as applied to generative AI. The case was not contested and swiftly reached a settlement resulting in a consent order restraining Dudesy from using Carlin's image, voice, or likeness.³⁶ A similar action was threatened by actress Scarlett Johansson, who accused OpenAI of using a voice for ChatGPT that was too similar to her own.³⁷ Again, this resulted in a swift capitulation by the putative defendant, leaving no clear answer as to how this tort may apply.

In at least one case, however, a complaint alleging that the use of "deepfake" technology violated the right of publicity survived a motion to dismiss.³⁸ In *Young v. Neocortex*, a self-identified "internet personality" sued a smartphone application that allegedly used "deepfake" software to place celebrities in different fanciful situations.³⁹ In denying the defendants' motion to strike, the court noted that the right of publicity does not fall within the subject matter of copyright and so is not preempted by the Copyright Act.⁴⁰ The Central District of California explained that unlike copyright, which protects specific works, the right to publicity encompasses the use of a person's image as a marketing tool or to imply association or endorsement, regardless of whether any particular copyrighted work was infringed.⁴¹

For those who claim their personal likeness or identity are being spoofed by generative AI, then, privacy or publicity rights may be an effective way to stop at least some potential harm. Claims based on publicity rights may avoid some of the more complex or novel questions about generative AI because they are agnostic as to the particular tool being used to infringe on privacy or publicity. Conceptually, they treat generative AI as nothing more than a tool. The focus is on the product of the generative AI software, its commercial use, and its impact on the economic value of another rather than on the functioning of the model itself.

Courts Appear Skeptical About the Theory of Generative AI as Mere Compression

Some plaintiffs are pushing the theory that a generative AI model represents a “compression” of the training data. Essentially, these plaintiffs argue that generative AI models “memorize” their training data and somehow “compress” it into models astronomically smaller than the original data set. Thus, the model is itself an infringing derivative work since it contains copies of copyrighted material, or so the argument goes.

Courts appear skeptical. In *Kadrey v. Meta Platforms, Inc.*, the plaintiff argued that “the infringing works are being retained inside the model” and thus the model itself is a direct infringement of copyright⁴² and every output from the model is an infringing derivative work.⁴³ The court expressed confusion, saying that “if you put the LLaMA⁴⁴ language model next to Sarah Silverman’s book you would say they’re similar” and that this “makes my head explode when I try to . . . understand that.”⁴⁵ The court then dismissed this aspect of the complaint, saying that the idea that the LLM itself is an infringing work “is nonsensical” and “there is no way to understand the LLaMA models themselves as a recasting or adaption of any of the plaintiff’s books.”⁴⁶

In *Andersen v. Stability AI, Ltd.*, one of the earliest cases brought by the same lawyers who filed *Kadrey*, the Northern District of California dismissed much of the original complaint, giving leave to amend.⁴⁷ The California court, as explained in the author’s 2023 series, was skeptical that five billion images could possibly be compressed to the size of an AI model⁴⁸ and noted that “the diffusion process involves not copying of images, but instead the application of mathematical equations and algorithms to capture concepts from the Training Images.”⁴⁹ It granted the plaintiff leave to amend to clarify its compression theory.⁵⁰

The *Andersen* plaintiffs complied, filing an amended complaint that provided new details about their compression theory.⁵¹ It proposed that image-generating AI takes preexisting images and then interpolates between them to create a new image.⁵² This assumes, however, that the AI model starts with a copyrighted work

and then selects which of them to interpolate each time a prompt is provided, which may not be an accurate description of the software. In a sense, this theory simply begs the question of whether compressed copies of copyrighted works lurk in the finished trained model.

To address this question, the *Andersen* plaintiffs cited to a study by Nicholas Carlini⁵³ concerning whether and to what extent generative AI memorizes specific training data.⁵⁴ The plaintiffs argued that because some images from the training data could be largely recreated by the model, the model must be compressing its training data.⁵⁵ The defendants argued that this was an overstatement because the study only showed that if an image was duplicated many times in the training data, it was sometimes possible to prompt the model to recreate something strikingly similar. The study started by looking for often-duplicated training examples.⁵⁶ It selected the most-duplicated examples and then generated 175 million images using prompts.⁵⁷ Ultimately, it found a total of 109 copies or near copies of the training data.⁵⁸ As the defendants in *Andersen* pointed out in their renewed motion to dismiss, this suggests that a generative AI model has a copy rate of about “one-in-a-million,” and then only if the target is overrepresented in the training data.⁵⁹

The court decided that the argument was sufficient to survive a motion to dismiss and that the truth “will be tested at a later date.”⁶⁰ The court differentiated the copying potential of generative AI from that of VCRs by noting that there was at least allegedly some direct evidence of the model’s designers’ specific intent to facilitate infringement.⁶¹ The court found that the disputed allegations about the plaintiffs’ compression theory were “sufficient to allow the direct infringement claims to proceed” but whether they are sufficient to support the claims “will be addressed at summary judgment.”⁶²

The New York Times, in its own lawsuit against Microsoft and OpenAI, supported its complaint with a similar argument.⁶³ The *Times*’ pleading included many examples of CoPilot- or ChatGPT-generated output that was identical or very similar to *Times* articles.⁶⁴ While these examples are striking, they do not reveal all of the prompts that generated the output or what

fraction of the prompts resulted in alleged copied output.⁶⁵ Whether and how frequently an LLM will produce memorized data probably depends on the particular LLM in question. When dealing with an “unaligned” LLM that has not been trained or pre-prompted by the developer to try to block or avoid producing memorized work, some studies are able to recover training data somewhere between 0.000001% and 0.852% of the time, depending on the model and methodology.⁶⁶

The likelihood of infringing output may matter. It may fall to the finder of fact in cases like these to analyze whether the memorized works sprang up regularly and unbidden, or whether they had to be carefully coaxed out of the model through precise or extensive prompt engineering. After all, the Supreme Court has held that “copying equipment, like the sale of other articles of commerce, does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes.”⁶⁷ If generative AI is aligned such that it can only produce infringing output when the user engages in prompt engineering to force it to do so, then there seems very little difference between this technology and accepted devices like cameras or copy machines.

Nevertheless, LLMs seemingly can, to some greater or lesser degree depending on the particular work, produce output closely resembling the training data.⁶⁸ The plaintiffs pointing this out may be more focused not on the fact that this is possible, but on the idea that because it is possible there must be “compressed copies” of the work within the model. This is a factual question, as the court in *Andersen* noted, but it seems like a difficult one. As mentioned in prior articles, the sheer difference in size between the training set and the models should give rise to some initial skepticism. Worse, the theories being pushed appear to be pointing toward a world where only the most famous and well-known artistic works earn robust protection against generative AI and smaller creators have none.

It is certainly true that reducing a copyrighted work to another medium does not alone immunize the copying.⁶⁹ Yet unlike digital copies that encode the copyrighted information with

fidelity limited only by the size of the new file, LLMs do something fundamentally different. A model is a set of relationships between prompts and output, not a stored copy of the training data. It is a map from input to output. While the roads on that map may have been trained on copyrighted data, the model's sensitivity to input is arguably more akin to a photocopier or VHS recording device than to a simple digital copy.

From a technical perspective, it is difficult to say that an LLM stores copies of anything in particular. Rather, the more frequently the model was trained on a specific work, the more the mathematical shape of the model maps words to aspects of that work. Each part of the training data tugs the web of algebra making up the model a little bit in its own direction. Famous works of art like the Mona Lisa, well-respected literature like *The Great Gatsby*, or often-quoted news sources like *The New York Times* likely appear many times in the training data and so pull the model harder, making parts of the model closely fit that work.⁷⁰ But the web of mathematics making up the model is also being pulled by everything else it has ingested. Therefore, only in extreme cases with the most well-represented works would a near-perfect copy appear to be possible. For most works, the relationships between a prompt and a particular training example would be warped so severely by everything else in the training data that perhaps only the vaguest impression of the original remains.

The Struggle to Align Generative AI With Copyright Laws or Break That Alignment

Operators of generative AI have not been insensitive to the risk that models may be used to create copyrighted work and have been taking steps to prevent such use. To begin with, commercial models like ChatGPT have been subjected to reinforcement learning through human feedback.⁷¹ After the model is trained on its dataset and perhaps has memorized some of it, the model is then subjected to additional training where humans provide feedback on the quality of its responses.⁷² This, like any other training, warps and changes the map of prompts to outputs and should make it more difficult for a model to produce copyrighted

work. Additionally, a model can probably be trained to specifically avoid memorization, as found by some researchers.⁷³

Copyright alignment does not just have to come from the model itself, though. Commercial applications like ChatGPT can supplement the user's prompt before it is sent to the model.⁷⁴ Operators of generative AI describe these as "guardrails" against infringement.⁷⁵ ChatGPT appears to provide some preliminary text that is provided to the model before the user's own prompt. As of July 2024, this includes instructions to "not name or directly/indirectly mention or describe copyrighted characters" and to "not discuss copyright policies in responses."⁷⁶ On top of this, ChatGPT uses software to monitor the activity of its model and flag certain kinds of problematic behavior.⁷⁷ If the model begins to produce part of a memorized copyrighted work, it is possible that the response will be terminated with a message reading, "This content may violate our Terms of Use or usage policies."⁷⁸

Paradoxically, while the defendant operators of generative AI appear to be taking various steps to prevent the generation of improper material, some of the plaintiffs suing them are working hard to find ways to coax copyrighted material out of the model. In *Walters v. OpenAI*, a journalist sued OpenAI in Georgia state court, asserting defamation when ChatGPT falsely reported that he was a defendant in a different lawsuit.⁷⁹ OpenAI claimed, however, that the user who generated the false information had to put in significant effort to generate the false information.⁸⁰ It further claimed that once the false information was generated, the user contacted the plaintiff journalist and then the two tried, and failed, to generate the same false report a second time.⁸¹ OpenAI argued that the plaintiffs were therefore eager participants in the conduct they now claimed was infringing.⁸²

In *Andersen*, Midjourney, one of the defendants, explained that the plaintiffs were able to obtain some images that they claim are similar to copyrighted work by feeding the model new copyrighted images as part of the prompts.⁸³ This is possible because Midjourney's generative AI model allows users to provide both an image and a text prompt, both of which are considered

when the model produces output.⁸⁴ Analogizing to a photocopier, Midjourney complained that the *Andersen* plaintiffs were only able to make copies by "feeding their own images into the tool."⁸⁵ The plaintiffs did not disagree that they did so.⁸⁶ Instead, they complained that this photocopier-like effect is itself a problem, making the model into a "copyright-laundering facility, designed to produce low-cost knockoffs . . ."⁸⁷

In *Times*, Microsoft alleged that the only reason the plaintiff was able to get some verbatim *Times* quotes out of the model was because it had essentially "hack[ed]" the generative AI model.⁸⁸ OpenAI claimed that it "took [the *Times*] tens of thousands of attempts to generate the highly anomalous results" that form the basis of the complaint, and to do so the *Times* had to "exploit[] a bug . . . by using deceptive prompts that blatantly violate OpenAI's terms of use."⁸⁹ The *Times* responded by arguing that it is not relevant how they got the model to produce the allegedly infringing results, only that they succeeded at all.⁹⁰

Evolving Business, Policy, and Legal Issues

Regardless of what the courts rule in the pending cases against generative AI, the technology is probably here to stay. And, as shown below, it is having significant effects on businesses, consumers, workers, and lawyers in the real world.

Unpredictability Leading to Problems for Business Use

Because generative AI is inherently unpredictable, those using it in a business context must be wary. One of the key concerns outlined in the second article in the author's 2023 series was that any business using generative AI as a customer-facing chatbot risks getting into trouble under contract, promissory estoppel, or similar theories.⁹¹ And, indeed, such problems have emerged. A Chevy dealership implemented a chatbot apparently powered by ChatGPT.⁹² Users quickly realized that they could manipulate the chatbot to produce unintended behavior such as writing python script, and another person was able to get the chatbot to agree to sell him an expensive truck for one dollar.⁹³

Luckily for the dealership, there does not appear to have been legal action to enforce the chatbot's promise. In a different case, however, a customer brought formal action against an airline based on his interaction with the airline's chatbot.⁹⁴ The customer asked Air Canada's chatbot how bereavement fare reductions worked and was told he could apply for them retroactively.⁹⁵ He relied on this and bought tickets, but Air Canada claimed the chatbot was in error and refused to reduce the fee.⁹⁶ The Civil Resolution Tribunal (CRT)⁹⁷ held that Air Canada was liable of negligent misrepresentation because it "did not take reasonable care to ensure its chatbot was accurate."⁹⁸

The holding in *Air Canada* may have less to do with the law of generative AI and more to do with inattentive lawyering. The CRT noted that Air Canada claimed its terms and conditions made the actual policy clear, but inexplicably "did not provide a copy of the relevant portion" of those conditions to the tribunal.⁹⁹ Thus, it remains to be seen whether robust disclaimers, clickwrap contracts, or similar methods will insulate a business from liability caused by a misbehaving chatbot. As discussed below, Colorado businesses should implement such agreements before introducing a customer facing chatbot if they wish to prevent discriminatory or harmful conduct.¹⁰⁰

Larger Policy Issues

While individual lawsuits may rise and fall based on the specifics of the technology or the elements of fair use, many of the plaintiffs suing generative AI companies are concerned about a larger shift in society. Characterized by some of the defendants as "anti-AI polemic," several of the complaints contain sweeping declarations of the feared harm to the creators' professions as a whole or even to humanity itself.¹⁰¹ As mentioned in earlier articles, the courts are probably not well suited for these expansive policy arguments.

The recent dismissal of *Cousart v. OpenAI LP* is an example.¹⁰² This lawsuit purported to include anyone who may have written information online that was used to train a model, anyone who had ever used ChatGPT, and anyone who had ever used a Microsoft

product.¹⁰³ It opened with a quote threatening the end of human civilization and continued with similar expansive language throughout.¹⁰⁴ The causes of action concerned the alleged theft or misappropriation of data posted online by individuals.

The Northern District of California dismissed the lawsuit without ever reaching the specific claims.¹⁰⁵ Invoking Rule 8(a), the court dismissed with leave to amend because "the complaint is not only excessive in length, but also contains swaths of unnecessary and distracting allegations making it nearly impossible to determine the adequacy of the plaintiffs' legal claims."¹⁰⁶ Among them were "five pages on how various political leaders and European governments have reacted to recent advancements in AI technology" and "rhetoric and policy grievances that are not suitable for resolution by federal courts" such as comparing the risk of AI to that of nuclear weapons.¹⁰⁷ The court gave the plaintiffs 21 days to get "the mud . . . off the walls of the complaint"¹⁰⁸ and try again.

Though their concerns might be better directed to the legislature, authors and artists have reason to be concerned. When generative AI was new, the only data on how prevalent the new technology would be in the workplace was speculative, much of it based on surveys.¹⁰⁹ Now that generative AI has been with us for over a year, empirical data suggests it is causing actual harm to individuals in the job market.

Two studies on websites that connect freelancers with clients have shown declines in the amount of work available for freelancers.¹¹⁰ The first study examined Upwork, "one the largest online labor markets in the world."¹¹¹ The study focused on 92,457 workers and 519,577 individual freelance gigs from January 2022 through April 2023.¹¹² It found significant negative effects coinciding with the release of ChatGPT during that time consisting of a "persistent and growing decrease in both the monthly number of jobs . . . and monthly compensation . . ."¹¹³ The study saw total monthly compensation in the most affected occupations dropping by 5.2%.¹¹⁴ These effects were more significant for projects affected by image-generating AI.¹¹⁵ This may corroborate anecdotal reports from artists saying that entry

level and freelance jobs in film, TV, and gaming are fewer and farther between.¹¹⁶

The second study examined evidence from both the introduction of ChatGPT and an earlier introduction of Google's neural network for language translation.¹¹⁷ First, this study examined 28,158 translations on Amazon Mechanical Turk from January 2016 through May 2017, spanning the time before and after Google introduced a neural-network-based translation service.¹¹⁸ It found a 13% to 20% drop for regular, analytical translations.¹¹⁹ This suggested a \$352,000 total loss of earnings to humans.¹²⁰ Looking to ChatGPT, the study next examined the number of questions and answers posted in Stack Overflow, a question and answer site used by programmers.¹²¹ It found a significant drop in the number of questions asked after the release of ChatGPT (that then partially rebounded), and a steady drop in the number of answers posted.¹²² The implication is that people are using ChatGPT or similar systems to help them with their programming questions, satisfying some of the need that would otherwise be filled by people. This conclusion is supported by other studies on similar question and answer sites in Chinese or Russian that have less access to ChatGPT.¹²³

Stack Overflow is not only seeing fewer users due to generative AI but also laying off workers to replace them with AI.¹²⁴ It is not the only company blaming layoffs on generative AI. Duolingo, the language learning app, laid off 10% of its employees as it moved to rely on AI.¹²⁵ Dropbox reportedly is cutting 16% of its staff, citing the use of AI.¹²⁶ While it is unclear exactly how many jobs have been lost to generative AI, some commentators think the number is understated by companies who do not want to generate bad press.¹²⁷

Some studies appear to verify that using generative AI can make workers more productive. One study finds that those who use generative AI in their own work experience an average 50% increase in productivity and favorability.¹²⁸ Another study of 444 experienced, college-educated professionals found that the time taken to complete a 20- to 30-minute writing assignment in their field dropped 37% among those using ChatGPT.¹²⁹ While increased

productivity is positive for production, it may also contribute to more pressure on wages and jobs as greater supply of productive work drives down demand for workers.

Legislation to Address Bias Concerns With Generative AI

Generative AI can reflect whatever biases or prejudices exist in its training data. Knowing this, Colorado has enacted a first-in-the-nation law governing bias in generative AI.¹³⁰ Previously, Colorado statute prohibited discrimination by insurance companies' algorithms,¹³¹ but the new law greatly expands the scope of the kind of industries that must guard against this problem.

The new law amends the Consumer Protection Act to mitigate "algorithmic discrimination," defined as "unlawful differential treatment or impact that disfavors" people based on a protected class.¹³² Effective in February 2026, developers and deployers of "high risk artificial intelligence systems" must "use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination."¹³³ However, the law specifically does not create any private right of action and can only be enforced by the Colorado Attorney General.¹³⁴

A "high risk" system under this new law is any AI system that, when deployed, makes or is a substantial factor in making decisions involving "the provision or denial to any consumer of, or the cost or terms of" education, employment, financial or lending service, government service, health care, housing, insurance, or legal service.¹³⁵ It does not include, however, many older and commonly used forms of software such as anti-virus software, video games, cybersecurity, spell-checking, web hosting, or others.¹³⁶ It also appears to have an exception for customer-facing chatbots so long as they are for the purpose of providing information, making referrals or recommendations, or answering questions, and so long as they are subject to an accepted use policy prohibiting generating discriminatory or harmful content.¹³⁷

The new law requires deployers to provide various notices to consumers of the technology and its uses.¹³⁸ Developers have various reporting and disclosure requirements, both

to those deploying systems as well as to the state of Colorado.¹³⁹

Lawyers and Generative AI

It remains to be seen whether the practice of law will benefit or suffer from the use of generative AI. On the one hand, there may be no obvious reason why the legal profession would be immune from the same market pressures and risk of substitution that seem to be harming other creative workers. On the other hand, the possibility of improved efficiency is an attractive way for a law firm to improve its performance.

New American Bar Association Model Rule on Generative AI

If a firm is going to employ generative AI, it should do so mindfully. To begin with, there are ethical risks and pitfalls related to diligence, candor, confidentiality, and bias.¹⁴⁰ Following over a year of discussion in the legal profession on the topic, the American Bar Association (ABA) issued Formal Opinion 512 addressing the use of generative AI in July of 2024.¹⁴¹ The opinion generally reiterates the concerns with competence, confidentiality, candor toward the tribunal, and supervisory duties covered in the author's 2023 article series. It raises two additional concerns, however, regarding fees and client consent.

The ABA notes that if a lawyer incurs an expense to use a generative AI tool to accomplish a task in a lesser amount of time, the lawyer may only charge for their actual time incurred but probably can pass along the cost of the tool.¹⁴² Hopefully, few lawyers in Colorado need to be reminded not to bill for more time than they actually spent on a task. The ABA goes on, however, to note that a law firm must also take care to distinguish overhead expenses (e.g., costs for generative AI embedded in word processing programs) from client-specific expenses (e.g., costs for a tool to compress the review of voluminous contracts), and typically only charge the client for the latter.¹⁴³

As to client consent, the ABA correctly points out that even a private model, trained and contained entirely within a law firm, is

not immune to confidentiality concerns.¹⁴⁴ A well-funded firm may develop a private model that is trained on the firm's own confidential or privileged information in order to become better at assisting counsel. If this is done, there is some risk that the model will ingest information related to one client and restate it in connection with work for an entirely different client. So, the firm must first obtain informed consent from its clients about the risks and benefits of the technology because of the risk of disclosure between lawyers working for different clients within the firm.¹⁴⁵

Additional Examples of Misconduct Involving Generative AI

In the last year or so, a handful of attorneys have found themselves in ethical trouble for failing to properly supervise generative AI.¹⁴⁶ In each case, the ultimate issue was that the lawyer did not check the hallucinated citations produced by the LLM. The reasons why this happened varied. The New York lawyers blamed lack of understanding of the new technology.¹⁴⁷ The Colorado lawyer pointed to inexperience.¹⁴⁸ Michael Cohen's lawyers explained that numerous lawyers had been involved in the drafting, and it had apparently not been clear who was supposed to perform a final cite check.¹⁴⁹ In the Massachusetts case, the lawyer failed to double-check work product of an intern.¹⁵⁰ Finally, in the Second Circuit decision, the lawyer appears to have missed other deadlines and requirements, with the hallucinated citations being just one of many errors.¹⁵¹ Despite the different sources of the error, the cure in each case is the same: remembering that generative AI is a secondary source at best, and citations must be checked prior to filing anything with the court.


It is easy for a lawyer to check citations on their own work product. But what of the confidentiality and other issues posed by generative AI in the hands of other firm employees? In addition to the models like ChatGPT that can be reached from any web browser, automatic updates to software are adding AI capabilities to software commonly used by lawyers, like Windows and Adobe. For example, Microsoft Windows now incorporates "CoPilot," an LLM,

into Windows directly.¹⁵² CoPilot appears with an automatic Windows update and allows users to converse with a chatbot assistant that can do things like search the web, produce images, and help with a variety of other computer tasks.¹⁵³ Adobe Acrobat, used by many firms to read and modify PDF documents, now also incorporates AI features such as document summary.¹⁵⁴

Thus, whether a firm uses generative AI or not, it should adopt policies governing the use of generative AI and educate employees on those policies before it becomes a problem. Among other things, a firm should carefully scrutinize the privacy policy of any software using generative AI and make sure its data is being kept confidential and not used to train the model.¹⁵⁵

Similarly, any use of generative AI should be checked to see whether it triggers the application of Colorado artificial intelligence law, which applies to a decision with a material impact on “the provision or denial” of a “legal service.”¹⁵⁶ Read literally, this might mean that any law firm using any form of AI in the client intake process, such as asking ChatGPT to help provide information about a potential client, may be required to provide various deployer-specific disclosures.

Conclusion

The uses for and law surrounding generative AI will continue to evolve, but at least one thing seems clear: the technology is here to stay. It seems likely that the technology will weather its legal challenges in some form or another, and so no profession, not even lawyers, will escape unchanged. Practitioners are well advised to stay abreast of the evolving law in this area and should set aside time to explore and gain a working knowledge of the technology.¹⁵⁷ But caution is paramount. As Justice Maria Berkenkotter recently explained, lawyers should aim to understand generative AI at least to the same degree that they understand how to safely operate their car.¹⁵⁸ Until we have a new statute or appellate level case telling us how often to change our generative AI oil and brake pads, it is up to us in the community to seek out information and develop responsible uses and safeguards. 



Colin E. Moriarty practices with Underhill Law, P.C. in Greenwood Village. Focusing on business and commercial litigation and arbitration, he has litigated business disputes, construction and fabrication defect claims, employment discrimination lawsuits, subcontractor litigation, state RICO fraud lawsuits, civil theft disputes, insurance appraisal and adjustment disputes, and other lawsuits involving complex commercial and construction matters—colin@underhilllaw.com.

Coordinating Editors: K Kalan, kmkalan@yahoo.com; William F. Vobach, bill@vobachiplaw.com

NOTES

1. Moriarty, “The Legal Challenges of Generative AI—Part 1: Skynet and HAL Walk Into a Courtroom,” 52 *Colo. Law.* 40 (July/Aug. 2023), <https://cl.cobar.org/features/the-legal-challenges-of-generative-ai-part-1>; Moriarty, “The Legality of Generative AI—Part 2: I’m sorry, User. I’m afraid I can’t do that.,” 52 *Colo. Law.* 30 (Sept. 2023), <https://cl.cobar.org/features/the-legality-of-generative-ai-part-2>; Moriarty, “The Legal Ethics of Generative AI—Part 3: A robot may not injure a lawyer, or, through inaction, allow a lawyer to come to harm.,” 52 *Colo. Law.* 30 (Oct. 2023), <https://cl.cobar.org/features/the-legal-ethics-of-generative-ai-part-3>.
2. *UMG Recordings, Inc. v. Uncharted Lab’s, Inc.*, No. 24-04777 (S.D.N.Y. June 24, 2024); *UMG Recordings, Inc. v. Suno, Inc.*, No. 24-11611 (D.Mass. June 24, 2024); *N.Y. Times v. Microsoft Corp.*, No. 23-cv-11195 (S.D.N.Y. Dec. 27, 2023).
3. *N.Y. Times*, No. 23-cv-11195; *Basebanes v. Microsoft*, No. 24-cv-00084 (S.D.N.Y. Jan. 5, 2024); *Sancton v. OpenAI, Inc.*, No. 1:23-cv-10211 (S.D.N.Y. Nov. 21, 2023); *Concord Music Grp. v. Anthropic*, No. 23-cv-01092 (M.D.Tenn. Oct. 18, 2023); *Huckabee v. Meta Platforms*, No. 23-cv-09152 (N.D.Cal. Oct. 17, 2023); *Authors Guild v. Open AI, Inc.*, No. 23-cv-08292 (S.D.N.Y. Sept. 19, 2023); *Chabon v. Open AI*, No. 23-cv-04625 (N.D.Cal. Sept. 8, 2023); *Kadrey v. Meta*, No. 23-cv-03417 (N.D.Cal. July 7, 2023); *Tremblay v. OpenAI, Inc.*, No. 23-cv-03223 (N.D.Cal. June 28, 2023).
4. *Flora v. Prisma Lab’s*, No. 23-cv-00680, 2023 U.S. Dist. LEXIS 138119 (N.D.Cal., Aug. 8, 2023); *Young v. Neocortex, Inc.*, 690 F.Supp. 3d 1091 (C.D.Cal. 2023); *P.M. v. OpenAI LP*, No. 23-cv-03199 (N.D.Cal. June 28, 2023); *Main Sequence, Ltd. v. Dudesy, LLC*, No. 2:24-cv-00711 (C.D.Cal. Jan. 25, 2024).
5. See *Levoy v. Alphabet Inc.*, No. 23-cv-03440 (N.D.Cal. July 11, 2023); *Cousart v. OpenAI LP*, No. 23-cv-04557 (N.D.Cal. Sept. 5, 2023).
6. Courts have previously held that compiling copyrighted material to create thumbnails or a searchable database is fair use. See *Authors Guild v. Google, Inc.*, 804 F.3d 202, 215 (2d Cir. 2015); *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1164–65 (9th Cir. 2007); *Kelley v. Arriba Soft Corp.*, 336 F.3d 811, 820–21 (9th Cir. 2003). They have also held that copying for the purpose of reverse engineering is fair use. See *Sony Comput. Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 602–03 (9th Cir. 2000); *Sega Enters. v. Accolade, Inc.*, 977 F.2d 1510, 1522 (9th Cir. 1992). It would seem that whether one views generative AI as a search tool or as reverse engineering the style or ideas behind the creative works they train on, fair use is likely to be a powerful argument, perhaps the key argument, in the legality of generative AI.
7. Order Granting in Part and Denying in Part Motions to Dismiss, *Andersen v. Stability AI Ltd.*, No. 23-cv-00201, 2024 U.S. Dist. LEXIS 143201 (N.D.Cal. Aug. 12, 2024).
8. *Doe v. GitHub, Inc.*, No. 22-cv-06823, 2024 U.S. LEXIS 11068 (N.D.Cal. Jan. 22, 2024).
9. *Andersen v. Stability AI Ltd.*, 700 F.Supp.3d 853, 875 (N.D.Cal. 2023).
10. 7 USC § 301(a).
11. *Thomson Reuters Enter. Ctr. v. Ross Intel. Inc.*, 694 F.Supp.3d 467 (D.Del. 2023). This case, while not dealing with generative AI, deals with such similar issues that litigants in the other cases are likely watching it closely. It was scheduled to go to trial in August 2024, but as of the date of this article it appears that trial has been postponed while the judge revisits summary judgment arguments.
12. *Id.* at 487.
13. *Id.*
14. *Id.* at 488.
15. *Id.* at 475.
16. *Id.* at 475.
17. *Id.* at 487–88.
18. *Id.*
19. *Doe v. GitHub, Inc.*, 672 F.Supp. 3d 837, 859 (N.D.Cal. 2023).
20. See 17 USC § 102(b). See also *Baker v. Selden*, 101 U.S. 99 (1880) (recently cited by *Google LLC v. Oracle Am. Inc.*, 593 U.S. 1, 47, (2021)); *Midler v. Ford Motor Co.*, 849 F.2d 460 (9th Cir. 1988) (applied in *Lewis v. Activision Blizzard, Inc.*, No. C12-1096, 2012 U.S. Dist. LEXIS 151739 (N.D.Cal. Oct. 22, 2012)); *Green v. Luby*, 177 F. 287 (Cir.Ct.S.D.N.Y. 1909); *Bloom & Hamlin v. Nixon*, 125 F. 977 (E.D.Pa. 1903).

21. See *Zacchini v. Scripps-Howard Broad. Co.*, 433 U.S. 562 (1977); *Haelan Lab'ys, Inc. v. Topps Chewing Gum*, 202 F.2d 866 (2d Cir. 1953). Described in Nimmer, "The Right of Publicity," 19 *Law and Contemp. Probs.* 203, 204 (Spring 1954), <https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=2595&context=lcp>.
22. See *Downing v. Abercrombie & Fitch*, 265 F.3d 994, 1004 (9th Cir. 2001).
23. Curren, "Digital Replicas: Harm Caused by Actors' Digital Twins and Hope Provided by the Right of Publicity," 102 *Tex. L. Rev.* 155 (Nov. 2023); McCaleb, "Paws Off My Profile: Protecting the Persona in a Modern Digital Age," 28 *Marq. Intell. Prop. & Innovation L. Rev.* 107 (Summer 2023).
24. *Donchez v. Coors Brewing Co.*, 392 F.3d 1211, 1220 (10th Cir. 2004) (rejecting a lawsuit by the creator of "Bob the Beerman" against Coors and noting in dicta that "the Colorado Supreme Court does not appear to have expressly recognized this tort").
25. *Joe Dickerson & Assocs. v. Dittmar*, 34 P.3d 995, 1001 (Colo. 2001).
26. *Id.* at 1002.
27. *Id.*
28. C.J.I. 28:1, n.3.
29. C.J.I. 28:4, n.2.
30. *Main Sequence, Ltd. v. Dudesy, LLC*, No. 24-cv-00711 at ¶¶ 38-46 (C.D.Cal. Jan. 25, 2024).
31. *Id.* at ¶ 47.
32. *Id.* at ¶ 45.
33. *Id.* at ¶¶ 41, 42, 46.
34. *Id.* at ¶¶ 47, 49, 54.
35. *Id.* at ¶¶ 80-91.
36. Stipulated Consent Judgment, *Dudesy*, No. 2:24-cv-00711 (C.D.Cal. June 18, 2024).
37. "Scarlett Johansson's Statement About Her Interactions With Sam Altman," *N.Y. Times* (May 20, 2024), <https://www.nytimes.com/2024/05/20/technology/scarlett-johansson-openai-statement.html>.
38. See *Young v. Neocortex, Inc.*, 690 F.Supp.3d 1091 (C.D.Cal. 2023).
39. *Id.*
40. *Id.* at 1102.
41. *Id.* at 1101-03.
42. Transcript of Proceedings 19, *Kadrey*, No. 23-cv-03417 (N.D.Cal. Nov. 17, 2023).
43. *Id.* at 9.
44. LLaMA stands for "Large Language Model Meta AI" and was one of the early open-source, downloadable LLMs released by Meta. See Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," arXiv:2302.13971 (Feb. 27, 2023), <https://arxiv.org/abs/2302.13971>.
45. *Id.* at 18.
46. Order Granting Motion to Dismiss, *Kadrey*, No. 23-cv-03417 (Nov. 20, 2023).
47. *Andersen*, 700 F.Supp.3d 853.
48. *Id.* at 865-66.
49. *Id.*
50. *Id.*
51. First Amended Complaint, *Id.*
52. *Id.* at 25-26.
53. *Id.* at 30.
54. Carlini et al., "Extracting Data from Diffusion Models," 1, published at the USENIX Security Symposium (Jan. 30, 2023), <https://arxiv.org/abs/2301.13188>.
55. Amended Complaint 34, *Andersen*, 700 F.Supp.3d 853.
56. Carlini, *supra* note 54 at 4.
57. *Id.* at 5.
58. *Id.* at 6.
59. DeviantArt, Inc.'s Motion to Dismiss 12, *Andersen*, 700 F.Supp.3d 853.
60. Order Granting in Part and Denying in Part Motions to Dismiss First Amended Complaint, *Andersen*, 2024 U.S. Dist. LEXIS 143204, *17.
61. *Id.* at *16-17.
62. *Id.* at *30-31.
63. *N.Y. Times*, No. 1:23-cv-11195.
64. *Id.* at 30-37.
65. *Id.*
66. Carlini et al., "Scalable Extraction of Training Data From (Production) Language Models," 1, 2, 8 (Nov. 28, 2023), <https://arxiv.org/pdf/2311.17035>.
67. *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417 (1984), superseded by 17 USCS § 1201(a)(2) (the Digital Millennium Copyright Act) with respect to avoiding copyright protection systems.
68. van den Burg and Williams, "On Memorization in Probabilistic Deep Generative Models" 2 (35th Conference on Neural Information Processing Systems, 2021), <https://proceedings.neurips.cc/paper/2021/file/eae15aabaa768ae4a5993a8a4f4fa6e4-Paper.pdf>.
69. *Meshwerks, Inc. v. Toyota Motor Sales U.S.A., Inc.*, 528 F.3d 1258, 1267 (10th Cir. 2008).
70. See van den Burg and Williams, *supra* note 68 at 2 (memorization is more likely when "the training data set contains a number of highly similar observations, such as duplicates of a particular work"). See also Yang, "Unveiling Memorization in Code Models," 8 (Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (April 14-20, 2024, Lisbon, Portugal), <https://doi.org/10.1145/3597503.3639074>.
71. Chaudhari et al., "RLHF Deciphered: A Critical Analysis of Reinforcement Learning From Human Feedback for LLMs," 3 <https://arxiv.org/pdf/2404.08555>.
72. In practice, the human actually trains an intermediate model, and the intermediate model, in turns, trains the underlying LLM. This solves the problem of humans not being fast enough to rate responses to train a model.
73. Kassem, "Mitigating Approximate Memorization in Language Models via Dissimilarity Learned Policy," 8 (May 2, 2023), <https://arxiv.org/pdf/2305.01550>.
74. ChatGPT is notably not just an LLM anymore. It also has the capacity to use Dall-E, an image-generating model, meaning users can obtain both text and image responses.
75. *Concord Music Grp., Inc. v. Anthropic PBC*, No. 3:23-cv-01092 (M.D.Tenn. Jan. 16, 2024).
76. See Conitzer and Leben, "How ChatGPT Has Been Prompted to Respect Safety, Fairness, and Copyright," Carnegie Mellon University (Feb. 26, 2024), <https://www.cmu.edu/tepper/news/stories/2024/february/chatgpt-safety-fairness-copyright.html>. See also Hetzschooldt, "Asking ChatGPT-4 About Its 'System Prompts', to Prevent Copyright Infringement," Substack (Feb. 9, 2024), <https://p4sc4l.substack.com/p/asking-chatgpt-4-about-its-system>.
77. Moderations, ChatGPT API Reference, <https://platform.openai.com/docs/api-reference/moderations/object>.
78. This was gleaned from the author's own experimenting with ChatGPT but is not difficult to reproduce.
79. Complaint, *Walters v. Open AI*, No. 23-A-04860-2 (Gwinnet Cnty.Super.Ct. June 9, 2023).
80. Motion to Dismiss, *Walters v. Open AI*, No. 1:23-cv-03122 (N.D.Ga. Nov. 1, 2023). The case was briefly removed to federal court but has since been remanded.
81. *Id.*
82. *Id.* The Georgia court denied the motion to dismiss without comment or analysis. Order, *Walters*, No. 23-A-04860-2 (Jan. 11, 2024). That might not be particularly surprising under Rule 12, though, since the motion appeared to rely heavily on allegations of facts outside the pleadings.
83. Midjourney's Motion to Dismiss 17, *Andersen v. Stability AI, Ltd.*, No. 3:23-cv-00201 (N.D.Cal. Feb. 8, 2024), Doc. 160.
84. Image Prompts, Midjourney, <https://docs.midjourney.com/docs/image-prompts>.
85. Midjourney's Reply in Support of Defendant Midjourney, Inc.'s Motion to Dismiss 12, *Andersen v. Stability AI, Ltd.*, No. 3:23-cv-00201 (N.D.Cal. Apr. 19, 2024), Doc. 184.
86. *Andersen* Response to Motion to Dismiss, *Id.* (Mar. 21, 2024), Doc. 176.
87. *Id.* at 18.
88. OpenAI's Memorandum of Law in Support of Motion to Dismiss, *N.Y. Times*, No. 1:23-cv-11195 (Feb. 26, 2024), Doc. 52.
89. *Id.*
90. "OpenAI's true grievance is not about how The Times conducted its investigation, but instead with what that investigated exposed . . ." Response to Motion to Dismiss 1, *N.Y. Times*, No. 1:23-cv-11195 (Mar. 11, 2024).
91. Moriarty, "The Legality of Generative AI—Part 2: I'm sorry, User. I'm afraid I can't do that.," *supra* note 1.
92. Notopoulos, "A Car Dealership Added an AI Chatbot to Its Cite. Then all Hell Broke Loose," *Business Insider* (Dec. 18, 2023), <https://www.businessinsider.com/car-dealership-chevrolet-chatbot-chatgpt-pranks-chevy-2023-12>.
93. *Id.*
94. *Moffat v. Air Can.*, 2024 BCCRT 149 (entered Feb. 14, 2024), <https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html>.

95. *Id.* at ¶ 2.
96. *Id.* at ¶ 4.
97. The CRT is “an independent, quasi-judicial tribunal” operating under the authority of a Canadian law. About the CRT, Civil Resolution Tribunal, <https://civilresolutionbc.ca/about-the-crt>.
98. *Moffat*, 2024 BCCRT at ¶ 28.
99. *Id.* at ¶ 31.
100. See CRS § 6-1-1701(9)(b)(R).
101. Defendant Google LLC’s Notice of Motion and Motion to Dismiss 1, *J.L. v. Alphabet Inc.*, No. 3:23-cv-03440 (N.D.Cal. Oct 16, 2023).
102. Order Granting Motions to Dismiss, *Cousart v. OpenAI LP*, No. 23-cv-04557, Doc. No. 78, (N.D.Cal. Sept. 5, 2023).
103. Complaint, *id.* at 64. (Feb. 27, 2024).
104. *Id.* at 1.
105. Order Granting Motions to Dismiss, *id.* (May 24, 2024).
106. *Id.*
107. *Id.*
108. *Id.*
109. *Id.*
110. Hui et al., “The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market,” SSRN (July 31, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4527336; Yilmaz et al., “AI-Driven Labor Substitution: Evidence From Google Translate and ChatGPT” (INSEAD working paper no. 2023/24/EFE, 2023), SSRN (Apr. 17, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4400516.
111. Hui, *supra* note 110 at 5.
112. *Id.* at 8.
113. *Id.* at 9, 10.
114. *Id.* at 9.
115. *Id.* at 12.
116. Carter, “Workers Are Worried About AI Taking Their Jobs. Artists Say it’s Already Happening,” *Business Insider* (Oct. 1, 2023), <https://www.businessinsider.com/ai-taking-jobs-fears-artists-say-already-happening-2023-10>.
117. Yilmaz, *supra* note 110 at 3.
118. *Id.* at 3.
119. *Id.* at 4.
120. *Id.* at 19. The more creative form of translation requiring adaption of meaning to cultural and emotional elements did not drop, however. *Id.* at 4. The authors hypothesized that this might suggest that the model struggled from a lack of “creative and cultural understanding” based on “personal and cultural experience, intuition and personal skills” that are difficult to formalize. *Id.* at 12.
121. *Id.* at 26.
122. *Id.* at 26–27.
123. Rio-Chanona et al., “Are Large Language Models a Threat to Digital Public Goods? Evidence from Activity on Stack Overflow” (July 14, 2023), <https://arxiv.org/pdf/2307.07367>.
124. Ghoshal, “Generative AI Forces Stack Overflow to Lay Off 28% of Its Workforce,” *InfoWorld* (Oct. 17, 2023), <https://www.infoworld.com/article/3708738/generative-ai-forces-stack-overflow-to-lay-off-28-of-its-workforce.html>.
125. Korn, “Duolingo Lays Off Staff as Language Learning App Shifts Toward AI,” *CNN* (Jan. 9, 2024), <https://edition.cnn.com/2024/01/09/tech/duolingo-layoffs-due-to-ai/index.html>.
126. Thorbecke, “AI is Already Linked to Layoffs in the Industry That Created It,” *CNN* (July 4, 2023), <https://www.cnn.com/2023/07/04/tech/ai-tech-layoffs/index.html>.
127. Constantz, “AI is Driving More Layoffs Than Companies Want to Admit,” *Yahoo! Finance* (Feb. 8, 2024), <https://finance.yahoo.com/news/ai-driving-more-layoffs-companies-174840542.html>.
128. Zhou and Lee, “Generative Artificial Intelligence, Human Creativity, and Art,” 3 *PNAS Nexus* 7 (Mar. 5, 2024) <https://doi.org/10.1093/pnasnexus/pgae052>.
129. Noy and Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence” (working paper, Mar. 2, 2023), https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf.
130. S.B. 24-205, https://leg.colorado.gov/sites/default/files/2024a_205_signed.pdf. California is close behind with its AI safety bill having passed the legislature. SB 1047, https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202320240SB1047. Governor Newsom has not yet signed it, and it is not clear that he will. Zeff, “Governor Newsom on California AI Bill SB 1047: ‘I can’t solve for everything,’” *Tech Crunch* (Sept. 17, 2024), <https://techcrunch.com/2024/09/17/governor-newsom-on-california-ai-bill-sb-1047-i-cant-solve-for-everything>. This bill differs from Colorado’s bill insofar as it seems to be concerned with the potential for large-scale, catastrophic risks as opposed to being narrowly focused on discrimination.
131. CRS § 10-3-1104.9.
132. CRS § 6-1-1701(1)(a). The introduced version of the bill also included a section regarding copyright issues, but this was removed from the final version as adopted. S.B. 24-205 (as introduced), https://leg.colorado.gov/sites/default/files/documents/2024A/bills/2024a_205_01.pdf.
133. CRS §§ 6-1-1702(a) and -1703(a).
134. CRS § 6-1-1706(6).
135. CRS § 6-1-1701(3) and (9)(a).
136. CRS § 6-1-1701(9)(b).
137. CRS § 6-1-1701(9)(b)(R).
138. CRS §§ 6-1-1703(4) and -1704.
139. CRS § 6-1-1702(2).
140. Moriarty, “The Legality of Generative AI—Part 3: A robot may not injure a lawyer or, through inaction, allow a lawyer to come to harm.,” *supra* note 1.
141. ABA, Formal Op. 512, *Generative Artificial Intelligence Tools* (July 29, 2024), https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf.
142. *Id.* at 13.
143. *Id.*
144. *Id.* at 7.
145. *Id.*
146. *United States v. Cohen*, No. 1:18-cr-00602 (S.D.N.Y. Mar. 20, 2024); *Smith v. Farwell*, No. 2282CV01197 (Mass.Sup.Ct. Feb. 12, 2024); *Park v. Kim*, No. 22-2507 (2d Cir. Jan 30, 2024); *People v. Crabill*, 23 PD J067 (Nov. 22, 2023); *Mata v. Avianca, Inc.*, No. 1:22-cv-01461 (S.D.N.Y. June 22, 2023).
147. *Avianca*, No. 1:22-cv-01461.
148. *Crabill*, 23 PD J067.
149. *Cohen*, No. 1:18-cr-00602.
150. *Farwell*, No. 2282CV01197.
151. *Kim*, No. 22-2507.
152. Microsoft, Copilot in Windows: Your Data and Privacy, <https://support.microsoft.com/en-us/windows/copilot-in-windows-your-data-and-privacy-3e265e82-fc76-4d0a-afc0-4a0de528b73a>; “Data, Privacy, and Security for Microsoft CoPilot for Microsoft 365,” Microsoft (June 20, 2024) <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy>.
153. Nield, “Here’s Everything You Can Do With Copilot, the Generative AI Assistant on Windows 11,” *Wired* (Nov. 5, 2024), <https://www.wired.com/story/microsoft-windows-11-copilot-generative-ai-assistant-tips>.
154. Adobe, Get AI-generated Overview and Summaries, <https://helpx.adobe.com/acrobat/using/ai-generated-summaries.html>.
155. See, e.g., Microsoft, “Data, Privacy, and Security for Microsoft CoPilot for Microsoft 365,” *supra* note 152; Adobe, Content Usage and Handling Practices, <https://helpx.adobe.com/acrobat/using/data-usage-and-handling.html>. Adobe in particular has been the subject of some recent controversy regarding its use of at least some of the data uploaded to some of its software for training its own models. Goldman, “Adobe Stock Creators Aren’t Happy With Firefly, the Company’s ‘Commercially Safe’ Gen AI Tool,” *Venture Beat* (June 20, 2023), <https://venturebeat.com/ai/adobe-stock-creators-arent-happy-with-firefly-the-companys-commercially-safe-gen-ai-tool>.
156. CRS § 6-1-1701(3) and (9)(a).
157. Karlik, “‘AI won’: Judges Caution Lawyers to Educate Themselves About Artificial Intelligence in the Law,” *Colorado Politics* (Mar. 14, 2024), https://www.coloradopolitics.com/courts/judges-caution-lawyers-to-educate-themselves-about-artificial-intelligence-in-law/article_c03b5428-d657-11ee-988e-571def4493f8.html.
158. *Id.*